

Climate Change & Longtermism

Author: John Halstead¹

Executive Summary

In this report, I will evaluate the scale of climate change from a *longtermist* point of view. [Longtermism](#) is the idea that influencing the long-term future, thousands of years into the future and beyond, is a key moral priority of our time.

In economics, longtermism is embodied by the idea that we should have a zero rate of 'pure time preference': we should not discount the welfare of future people merely because it is in the future. Economists who embrace a zero rate of pure time preference will tend to favour more aggressive climate policy than those who discount future benefits.

Climate change is a proof of concept of longtermism. Every time we drive, fly, or flick a light switch, each of us causes CO₂ to be released into the atmosphere. This changes the amount of CO₂ that is in the atmosphere for a very long time: unless we suck the CO₂ out of the atmosphere ourselves, concentrations only fall back to natural levels after hundreds of thousands of years. The chart below shows long-term CO₂ concentrations after different amounts of cumulative carbon emissions.

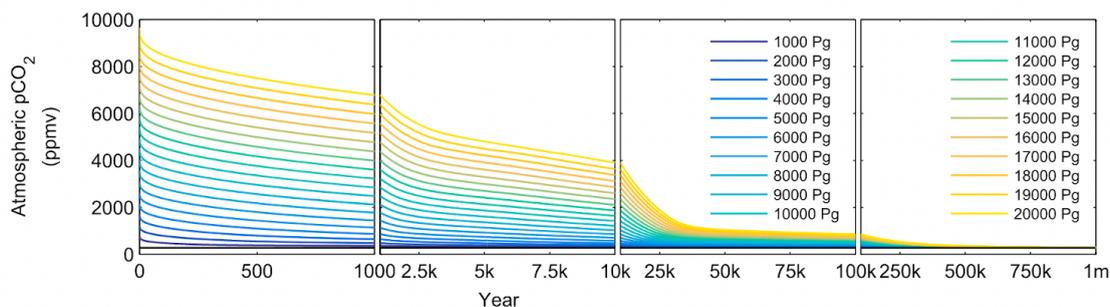


Figure 2. Atmospheric $p\text{CO}_2$ predicted by cGENIE for the pulse series scenarios (1000–20,000 Pg C). Preindustrial CO₂ concentrations are shown in black.

Source: N. S. Lord et al., 'An Impulse Response Function for the "Long Tail" of Excess Atmospheric CO₂ in an Earth System Model', *Global Biogeochemical Cycles* 30, no. 1 (2016): 2–17, <https://doi.org/10.1002/2014GB005074>.

Some of the ecological effects of climate change get worse over time. The clearest example of this is sea level rise. On current policy, the most likely sea level rise this century is 75cm. However, over 10,000 years, sea levels will rise by 10 metres. Over the long-term, the world will look very different.

¹ If you have any comments on this report or notice any errors, please email john.halstead309@gmail.com

From a longtermist point of view, it is especially important to avoid outcomes that could have persistent and significant effects. These include events like human extinction, societal collapse, a permanent negative change in human values, or prolonged economic stagnation. If we go extinct, then that would be the end of the human story, and there would be no future generations at all. If civilisation collapses permanently, then future generations will be left much worse off than they could have been, living lives full of suffering rather than ones of flourishing.

The anatomy of climate risk

The overall size of climate risk depends on the following factors:

1. Greenhouse gas emissions
2. The climate change we get from different levels of emissions
3. The impacts of different levels of climate change

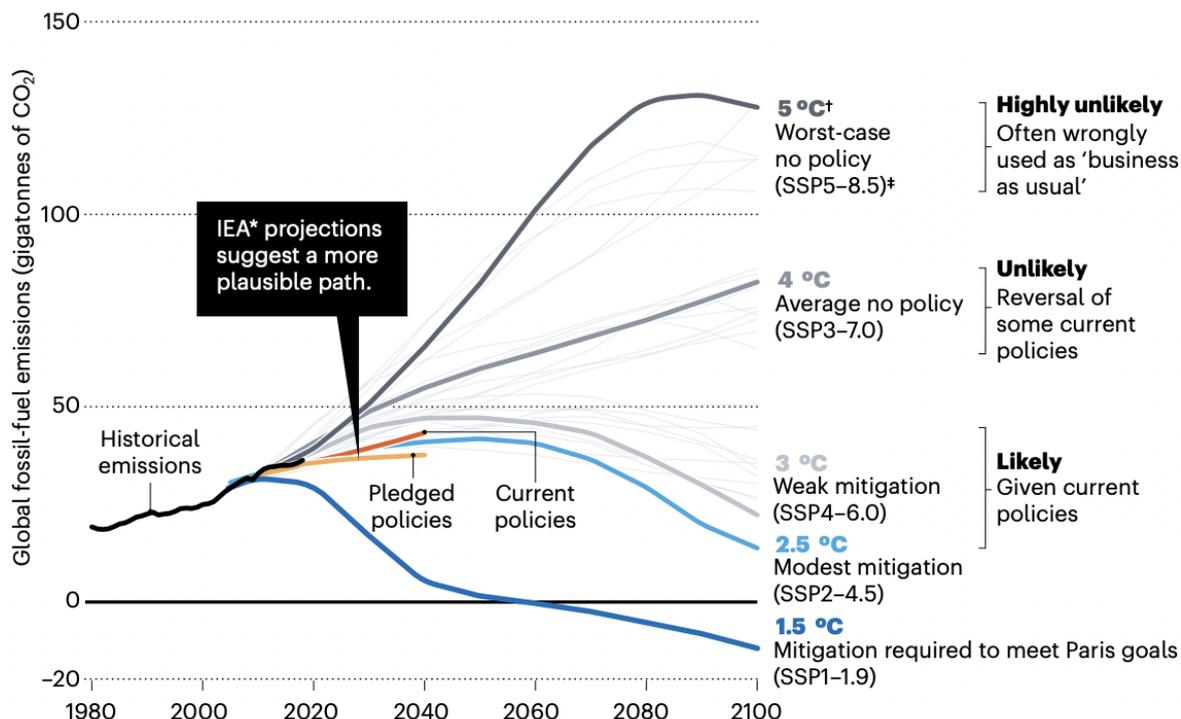
There is uncertainty about all three factors. The main findings of this report are as follows.

Emissions are likely to be lower than once thought

Due to recent progress on clean technology and climate policy, we look likely to avoid the worst-case emissions scenario, known in the literature as 'RCP8.5'. The most likely scenario on current policy is now the medium-low emissions pathway known as 'RCP4.5'. Moreover, climate policy is likely to strengthen in the future. For instance, as I was writing this report, the US Senate passed the Inflation Reduction Act, the most significant piece of climate legislation in American history.

POSSIBLE FUTURES

The Intergovernmental Panel on Climate Change (IPCC) uses scenarios called pathways to explore possible changes in future energy use, greenhouse-gas emissions and temperature. These depend on which policies are enacted, where and when. In the upcoming IPCC Sixth Assessment Report, the new pathways (SSPs) must not be misused as previous pathways (RCPs) were. Business-as-usual emissions are unlikely to result in the worst-case scenario. More-plausible trajectories make better baselines for the huge policy push needed to keep global temperature rise below 1.5 °C.



*The International Energy Agency (IEA) maps out different energy-policy and investment choices. Estimated emissions are shown for its Current Policies Scenario and for its Stated Policies Scenario (includes countries' current policy pledges and targets). To be comparable with scenarios for the Shared Socioeconomic Pathways (SSPs), IEA scenarios were modified to include constant non-fossil-fuel emissions from industry in 2018.

†Approximate global mean temperature rise by 2100 relative to pre-industrial levels.

*SSP5-8.5 replaces Representative Concentration Pathway (RCP) 8.5.

Source: Hausfather and Peters, [‘Emissions – the ‘business as usual’ story is misleading’](#), *Nature*, 2020.

Climate change is a great illustration of how society can make progress on a problem if enough people are motivated to solve it. This does not mean that climate change is solved, but there is significant momentum, and we are at least now moving in the right direction.

The amount of carbon we could burn in a worst-case scenario is also much lower than once thought. Some of the literature assumes that there are 5 or even 10 trillion tonnes of carbon remaining in fossil fuels, mostly in the form of coal. However, these estimates fail to recognise that not all fossil fuels resources are recoverable. Estimates of *recoverable* fossil fuels range from 1 to 3 trillion tonnes of carbon.

It is difficult to come up with plausible scenarios on which we burn all of the recoverable fossil fuels. Doing so would require (1) significant improvements in advanced coal extraction technology which is not part of the energy conversation today, but (2) a dramatic slowdown in progress in low carbon technologies that are already getting substantial policy support.

Warming is likely to be lower than once thought

Warming will likely be lower than once feared, in part because of lower emissions and in part because the scientific community has reduced uncertainty about climate sensitivity. Where once current policy seemed likely to imply 4°C of warming above pre-industrial levels, now the most likely level of warming is around 2.7°C, and the chance of 4°C is around 5%. Moreover, where once there seemed to be a >10% chance of 6°C on current policy, the risk now seems to be well below 1%.

On a worst-case scenario in which we burn all of the fossil fuels, the most likely level of warming is 7°C, and there is a 1 in 6 chance of more than 9.5°C.

Climate change will disproportionately harm the worst-off

The climate impacts literature suggests that climate change will impose disproportionate costs on countries at low latitude, which are disproportionately low- and middle-income and have done the least to contribute to climate change. People in Asia will have to deal with increasing flooding due to rising sea levels. Climate change will damage agricultural output, and cause droughts in countries reliant on rainfed agriculture. People in the tropics will face rising levels of heat stress. Fossil fuels also kill millions of people from air pollution in both poor and rich countries.

Many low- and middle-income countries have essentially never experienced sustained improvements in living standards, and a significant fraction may be left worse-off *than today* due to climate change. This undermines one common argument for discounting the future costs of climate change - that future generations will be richer and so better able to adapt to the effects of climate change.

We have a clear moral responsibility not to impose this harm, to reduce emissions, and to encourage economic development in poorer countries.

Average living standards will probably continue to rise

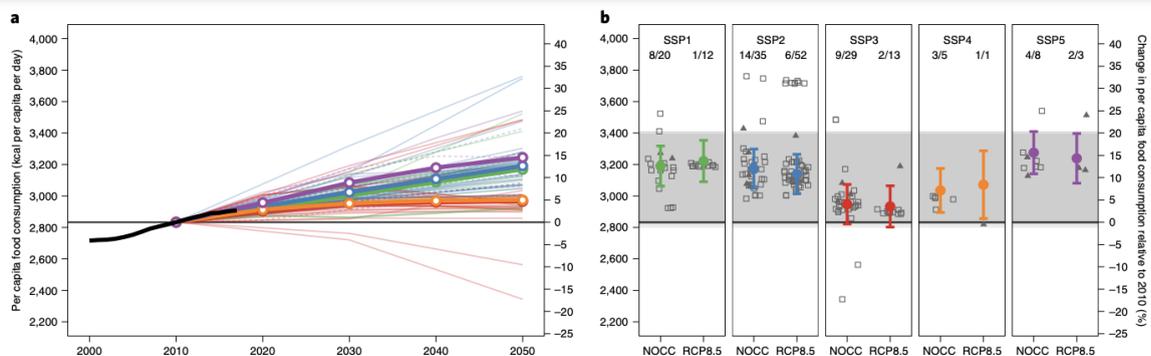
Climate-economy models confirm that the costs of climate change will fall disproportionately on poorer people, but most models also suggest that global average living standards in the future will be higher than today, on plausible levels of warming. Income per person looks set to increase by several hundred percent by the end of the century, notwithstanding the effects of climate change.

'Bottom-up' climate-economy models included in the IPCC's Sixth Assessment Report that add up the effects of climate impacts in different sectors and plug them into modern economic models suggest that warming of 4°C would do damage equivalent to reducing global GDP by around 5%. One recent model, Takakura et al (2019), includes the following impacts:²

² Jun'ya Takakura et al., 'Dependence of Economic Impacts of Climate Change on Anthropogenically Directed Pathways', *Nature Climate Change* 9, no. 10 (October 2019): 737–41, <https://doi.org/10.1038/s41558-019-0578-6>.

- Fluvial flooding
- Coastal inundation
- Agriculture
- Undernourishment
- Heat-related excess mortality
- Cooling/heating demand
- Occupational-health costs
- Hydroelectric generation capacity
- Thermal power generation capacity

For instance, in agriculture, the message from the climate impacts literature is that although climate change will damage food production, average food consumption per person will be higher than today, even for 4°C of warming, due to progress in agricultural productivity and technology. This is illustrated on the chart below from van Dijk et al (2021), which shows per capita food consumption on different socioeconomic pathways.



Source: Michiel van Dijk et al., 'A Meta-Analysis of Projected Global Food Demand and Population at Risk of Hunger for the Period 2010–2050', *Nature Food* 2, no. 7 (July 2021): 494–501, <https://doi.org/10.1038/s43016-021-00322-9>.

I have previously been critical of climate-economy models, but now believe they are more reliable than they once were. Until recently, a key determinant of aggregate impact assessments was how to model the effects of >4.4°C because the chance of that level of warming was so high. Estimates that models arrived at were unmotivated and arbitrary in part because the literature on the impacts of >4.4°C was sparse. However, warming of >4.4°C now seems increasingly unlikely (<1% given likely trends in policy), and there is a rich and voluminous literature on the impact of warming up to 4.4.°C. This makes recent bottom-up models more reliable.

However, even the best bottom-up climate-economy models underestimate the costs of climate change because they do not account for some important direct costs:

- They do not include tipping points
- They do not explicitly model the potential effects of climate change on economic growth and technological progress

It is unclear how much these factors would increase the overall direct costs of climate change; that is an important area of future research for climate economics. However, for

levels of warming that now seem plausible, these effects seem unlikely to be large enough to outweigh countervailing improvements in average living standards.

Bottom-up climate-economy models also do not account for indirect effects, such as conflict, which I discuss below.

'Top-down' climate-economy models try to directly measure the effects of climate change on aggregate economic output, and some of these find much higher impacts from climate change, on the order of a 25% reduction in GDP for 4°C warming. However, these results are highly model-dependent, rely on questionable econometric assumptions, and exclude several important climate impacts. In my view, the best bottom-up studies are a more reliable guide, notwithstanding their flaws.

Although average living standards are likely to continue to rise, we also need to consider the possibility of societal collapse for other reasons, such as a pandemic or nuclear war. If there were to be a major global catastrophe, then future living standards may not actually be higher than today. Future generations trying to rebuild society would have to do so in a less hospitable climate.

Some tipping points could have very bad effects

In my view, the most concerning tipping points highlighted in the literature are rapid cloud feedbacks, collapse of the Atlantic Meridional Overturning Circulation and collapse of the West Antarctic Ice Sheet.

Some models suggest that if CO₂ concentrations pass 1,200ppm (compared to 415ppm today), cloud feedbacks could cause 8°C of additional warming over the course of years to decades, on top of the 5°C we would already have experienced. The impacts of this sort of extreme warming have not been studied, but it seems plausible that hundreds of millions of people would die. Moreover, people would be stuck with an extreme greenhouse world for millennia. This would extend the 'time of perils': the period in which we have the technology to destroy ourselves, but lack the political institutions necessary to manage that technology. It would also make it much harder to recover from a civilisational collapse caused by something else (such as a pandemic or nuclear war). However, given progress on emissions, it is now difficult to come up with plausible scenarios on which CO₂ concentrations rise to 1,200ppm.

Collapse of the Atlantic Meridional Overturning Circulation would cause cooling and drying around the North Atlantic, and more importantly would probably weaken the Indian monsoons and the West African monsoons, with potentially dire humanitarian implications. For 4°C, models suggest that the chance of collapse is 1-5%, though they probably understate the risk.

There is deep uncertainty about potential sea level rise once warming passes 3°C. For higher levels of warming, there is a risk of non-linear tipping points, such as collapse of the West Antarctic Ice Sheet, which would cause sea levels to rise by around 5 metres over 100 years, which would probably cause flooding of numerous highly populated cities, especially in Asia.

Due to progress on emissions, these tipping points now look less likely than they did ten years ago, but their *expected* costs (impact weighted by probability) may still be large. Furthermore, our understanding of the climate system is imperfect, and there may be other damaging tipping points that we do not yet know about.

All this being said, contra some prominent research, the evidence from models and the paleoclimate (the deep climate history of the Earth) suggests that it is not the case that, once warming passes 2°C-4°C, runaway feedback loops will kick in that make the world uninhabitable.

Direct impacts fall well short of human extinction

Given progress in emissions, the risk of human extinction from the direct effects of climate change now seems extremely small. The most plausible route to human extinction is via runaway feedback loops. However, models and evidence from the paleoclimate suggest that it is impossible to trigger such runaway effects with fossil fuel burning. Models suggest that we could only trigger a runaway greenhouse if CO₂ concentrations pass 3,000ppm (at the very least), which is out of reach on revised estimates of recoverable fossil fuels.

Moreover, global average temperatures have been upwards of 17°C higher several times in the past without triggering runaway feedback loops that killed all life on Earth. Indeed, since the Cretaceous, 145 million years ago, periods of high temperatures and/or rapid warming have not been associated with ecological disaster. However, prior to the Cretaceous, climate change was linked to ecological disaster. In the report, I discuss the theory that this was because of ecological and geographical factors unique to the pre-Cretaceous period.

I construct several models of the direct extinction risk from climate change but struggle to get the risk above 1 in 100,000 over all time.

One argument that climate change could directly cause civilisational collapse is that it could be a contributing factor (along with deforestation, human predation, and pollution) to ecosystem collapse, which could in turn cause the collapse of global agriculture. I argue in the main report that this risk is minimal.

Indirect risks are under-researched but now seem fairly low

Because interstate war has become increasingly rare since the end of World War II, most of the literature on climate change and conflict has focused on the connection between climate and *civil* conflict: conflicts between a government and its citizens in which more than 25 people are killed.

Scholars in the field agree that, so far, climate-related factors have been a much weaker driver of civil conflict than other factors such as socioeconomic development and state capacity. However, there is strong disagreement in the field about how important climate change will be in the future. It is widely agreed that the risk of climate-induced conflict is

greatest in low- and middle-income countries, and that the most important mechanism is damage to agriculture.

The potential impact of climate change on the risk of *interstate*, rather than civil, war is potentially much more important but also much less studied. Among interstate conflicts, conflicts between the major powers pose by far the largest risk to humanity. This is because the major powers have far more destructive weaponry and have the capacity to alter the trajectory of humanity in other ways.

The most plausible way that climate change could affect the risk of interstate war is by causing agricultural disruption, which causes civil conflict, which in turn causes interstate conflict. Indeed, there is some evidence that countries embroiled in civil conflict are more likely to engage in military disputes with other countries.

It is difficult to see how climate change could be an important driver of some of the most potentially consequential conflicts this century - between the US and Russia, and the US and China. It is more plausible that climate change could play a larger role in driving conflict between India and Pakistan and also India and China. However, for plausible levels of warming, other drivers of this conflict seem much more important.

It is extremely difficult to provide reliable quantitative estimates of the risk of Great Power War caused by climate change. Nonetheless, I have built a model that attempts to put some numbers on the key considerations. I think this is valuable for several reasons. Firstly, it clarifies the cruxes of disagreements and allows focused discussion on those cruxes. Secondly, it allows us to prioritise different problems. If we do not quantify, we will still have judgments about how important different considerations are. Models make these considerations precise.

The downside of quantitative models is that they can cause false precision and anchor readers, even if the model is not good and has not been subject to scrutiny. Many of the considerations I have discussed are very difficult to quantify because there is essentially no literature on them.

With those caveats in my mind, my best guess estimate is that the indirect risk of existential catastrophe due to climate change is on the order of 1 in 100,000, and I struggle to get the risk above 1 in 1,000. Working directly on US-China, US-Russia, India-China, or India-Pakistan relations seems like a better way to reduce the risk of Great Power War than working on climate change.

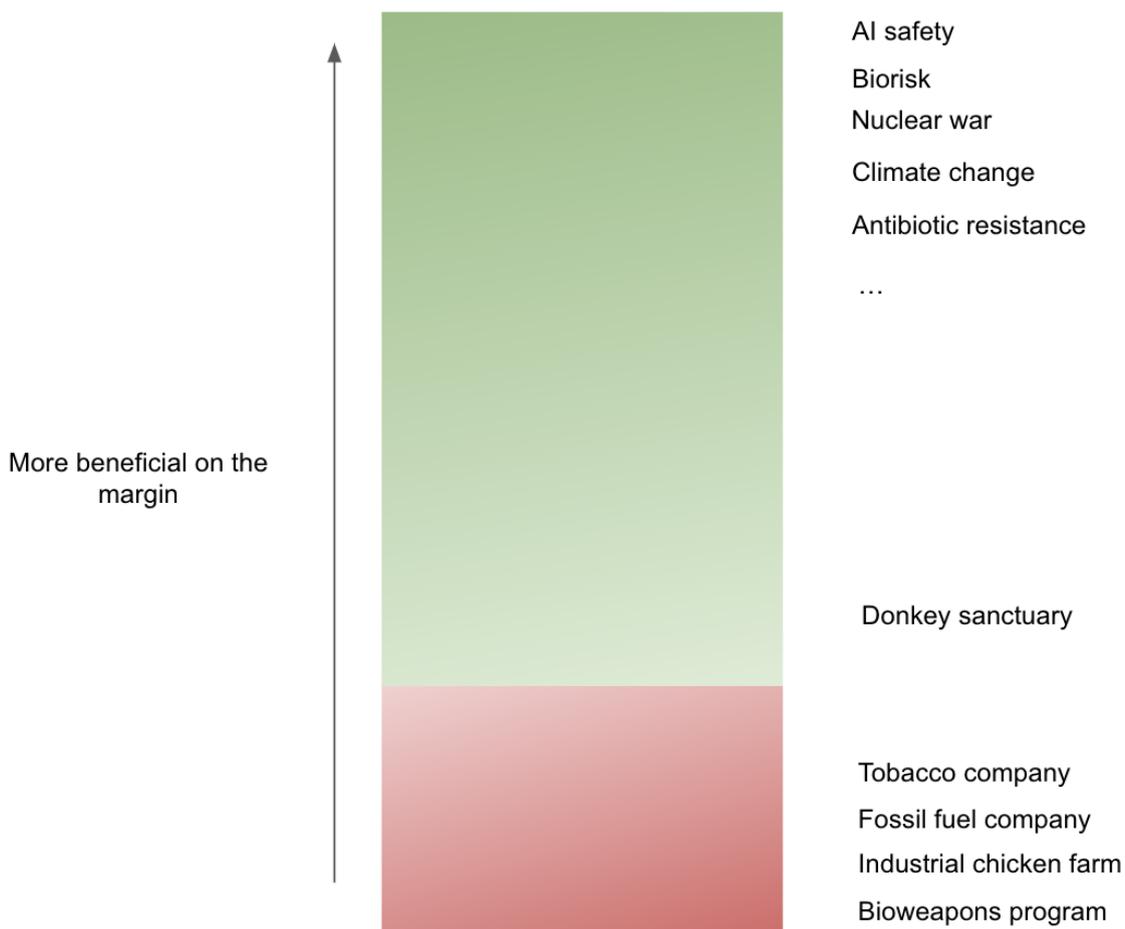
My personal thoughts on prioritising climate change relative to other problems

My primary goal in this report is to help people to answer the following question:

If your goal is to make the greatest possible positive impact on the world, what should you do with your time and money right now, given how the rest of society is spending its resources?

Crucially, this question is about what people should do *on the margin*. It is about what people should do given how society allocates its resources, not about how society as a whole should allocate its resources. Thus, when I say that working on some other problems, such as nuclear war or biosecurity, will have greater impact, this doesn't mean that society as a whole should spend nothing on climate change and everything on nuclear war and biosecurity. Rather, it is a claim about what we should do with our resources given how other resources are currently spent.

Moreover, the question I am trying to answer in this report is specifically about how to make the *greatest possible* impact on the world. This is the highest possible bar. In my view, climate change is one of the most important problems in the world, but other problems, including engineered viruses, advanced artificial intelligence and nuclear war, are more pressing on the margin because they are so neglected. One can visualise this in the following way. Green projects are beneficial on the margin, and red projects are harmful on the margin. Deeper green projects are more beneficial whereas deeper red projects are more harmful on the margin.



To emphasise, we should not confuse the claim that other problems are more pressing than climate change with the claim that climate change doesn't matter at all. I am glad that

climate change is a top priority for millions of young people and for many of the world's smartest scientists, and I would like governments and the private sector to spend more on climate change. I helped to set up the Founders Pledge Climate Change Fund ([donate here](#)), which has helped to move millions of dollars to effective climate change charities. The point is that I would like other global catastrophic risks to receive comparable attention, not that I would like climate change to receive less than it does today.

Imagine that only a few hundred people in the world thought that climate change is an important problem (rather than at least tens of millions), that philanthropists worldwide spent a few million dollars a year on climate (rather than \$10 billion), that society as a whole spent a million dollars on the problem (rather than \$1 trillion), and that the international institutions trying to tackle the problem either don't exist or have a similar budget to a McDonald's restaurant. How bad would climate change be? This is how bad things are for the other global catastrophic risks, and then some.

The final important piece of context is as follows: although I am taking a longtermist perspective in this report, my conclusions about the priority of climate change relative to other global catastrophic risks **are also true if you think only current generations matter**. In my view, the risks from AI, biorisk and nuclear war this century are much higher than commonly recognised.

- **AI:** Forecasters on the community forecasting platform Metaculus think that artificial intelligent systems that are better than humans at all relevant tasks [will be created in 2042](#). The most sophisticated attempt to forecast transformative AI is by Ajeya Cotra, a researcher at the Open Philanthropy Project and [her model now suggests](#) that it is most likely to be developed in 2040. A [2017 survey](#) of hundreds of leading AI researchers found that the median judgments implied that there is around a 4% chance of human extinction caused by AI before the end of the century.
- **Biorisk:** [Combined forecasts](#) on Metaculus imply that the chance of synthetic biology killing more than 10% of the world population by 2100 is around 7%. The implied chance of synthetic biology killing [more than 95% of the world population](#) before 2100 is around 0.7%.
- **Nuclear war:** Forecasters on the community forecasting platform Metaculus think that there is an 8% chance of [thermonuclear war by 2070](#).

These risks are not speculative possibilities, and the case for working on them is not contingent on ignoring the suffering of the current generation for the sake of a tiny probability of techno-catastrophe. I think it highly likely that my daughter will have to live through nuclear war, pandemics created by engineered viruses, and/or the emergence of transformative AI systems that will radically alter society. It is deeply unfortunate that few people acknowledge these problems, and that many people who are aware of them dismiss them as sci-fi fantasies without attempting to engage with the arguments, or grappling with the fact that *many people working in these fields* agree that the risks are large.

Although, I contend, my conclusions follow on both neartermist and longtermist perspectives, it is important to reiterate that, in my view, a longtermist ethical point of view is the correct one. I see no compelling arguments for ignoring the welfare of future generations, and an

ethical system that does ignore them is obviously difficult to square with concern about climate change.

While many people accept that the *direct* risks of climate change are lower than these other risks, some argue that the indirect effects of climate change may be large enough to make the total risk of climate change comparable. I do not think this is plausible. As discussed above, my rough models suggest that the total risk of climate change falls well short of the direct risk posed by the other global catastrophic risks. Moreover, the other risks *also* have indirect effects. As a rule, we should expect greater direct risks to have greater indirect effects. For instance, the indirect effects of trends in biotechnology seem to me much larger than the indirect effects of climate change. If biotechnology does democratise the creation of weapons of mass destruction, the indirect effects for the global economy and geopolitics are hard to fathom but seem enormous.

Overall, because other global catastrophic risks are so much more neglected than climate change, I think they are more pressing to work on, on the margin. Nonetheless, climate change remains one of the most important problems from a longtermist perspective. If progress stalls and emissions are much higher than we expect, then there is a non-negligible chance of highly damaging tipping points. Moreover, climate change is a stressor of political upheaval and conflict, which can in turn increase other global catastrophic risks. Finally, extreme climate change would make recovery from civilisational collapse more difficult.

Acknowledgments

I am very grateful to the following expert reviewers for their comments and criticism. Any mistakes are my own, and the views expressed here are my own.

- Matthew Huber, Professor, Dept. of Earth, Atmospheric and Planetary Sciences, Purdue University
- Dan Lunt, Professor of Climate Science, Bristol University
- Jochen Hinkel, Head of Department of Adaptation and Social Learning at the Global Climate Forum
- R. Daniel Bressler, PhD Candidate in Economics at Columbia
- Cullen Hendrix, Professor at the Korbel School of International Studies, University of Denver
- Andrew Watson, Royal Society Research Professor at the University of Exeter
- Peter Kareiva, Pritzker Distinguished Professor in Environment & Sustainability, UCLA
- Christina Schädel, Assistant Research Professor Center for Ecosystem Sciences and Society, Department of Biological Sciences, Northern Arizona University
- Joshua Horton, Research Director, Geoengineering, Keith Group, Harvard
- Laura Jackson, UK Met Office
- Keith Wiebe, Senior Research Fellow at the International Food Policy Research Institute
- Matthew Burgess, Assistant Professor, Department of Environmental Studies, University of Colorado Boulder
- David Denkenberger, Assistant Professor of Mechanical Engineering at University of Alaska Fairbanks
- Peter Watson, Senior Research Fellow and Proleptic Senior Lecturer, School of Geographical Sciences, Cabot Institute for the Environment, University of Bristol
- Goodwin Gibbins, Research Fellow, Future of Humanity Institute, University of Oxford
- Linus Blomqvist, Senior Fellow at Breakthrough Institute, PhD candidate in Environmental Economics and Science at UC Santa Barbara
- Luca Righetti, Research Fellow, Open Philanthropy Project
- Johannes Ackva, Climate research lead, Founders Pledge
- James Ozden, Extinction Rebellion

For helpful advice and discussion, thanks to Prof. Tapio Schneider, Dr. Roman Hoffman, Prof Colin Goldblatt, Dr. Zeke Hausfather, Prof. Paul Wignall, and Prof. Ramses Ramirez. Finally, I am grateful to my colleagues Max Daniel and Stephen Clare for their comments.

Table of Contents

[Executive Summary](#)

[The anatomy of climate risk](#)

[Emissions are likely to be lower than once thought](#)

[Warming is likely to be lower than once thought](#)

[Climate change will disproportionately harm the worst-off](#)

[Average living standards will probably continue to rise](#)

[Some tipping points could have very bad effects](#)

[Direct impacts fall well short of human extinction](#)

[Indirect risks are under-researched but now seem fairly low](#)

[My personal thoughts on prioritising climate change relative to other problems](#)

[Acknowledgments](#)

[Table of Contents](#)

[Introduction](#)

[1. How much will we emit?](#)

[1.1. Emissions so far](#)

[1.2. Future emissions](#)

[1.2.1. Climate progress](#)

[1.2.2. Assumptions about coal use](#)

[1.2.3. Assumptions about economic growth](#)

[1.3. Recent projections of emissions](#)

[1.3.1. International Energy Agency](#)

[1.3.2. Liu and Raftery \(2021\)](#)

[1.3.3. Climate Action Tracker](#)

[1.3.4. Integrated Assessment Models](#)

[1.3.5. Meinshausen et al \(2020\)](#)

[1.3.6. Summary](#)

[1.4. Will renewables and batteries solve climate change on their own?](#)

[1.5. Worst-case emissions scenarios](#)

[1.5.1. How high could fossil fuel emissions be?](#)

[Will we try and extract all of the coal?](#)

[Ultimately recoverable resources](#)

[1.5.2. Explosive economic growth](#)

[1.5.3. Indefinitely stalled decarbonisation](#)

[Excessive regulation](#)

[Low hanging fruit](#)

[Breakdown in international coordination](#)

[Overall assessment](#)

[1.6. Can we change course?](#)

[1.6.1. The physics](#)

[1.6.2. How hard is changing course?](#)
[Emissions reduction should be the priority](#)

[1.6.3. Crash direct air capture](#)

[1.6.4. How likely is crash direct air capture?](#)

[Observed willingness to pay to abate](#)

[How might willingness to pay change in the future?](#)

[2. How hot will it get?](#)

[2.1. Climate sensitivity metrics](#)

[2.1.1. Advantages and disadvantages of different metrics](#)

[Equilibrium metrics](#)

[Transient Climate Response to Cumulative Emissions](#)

[2.2. Estimates of climate sensitivity](#)

[2.2.1. Equilibrium Climate Sensitivity](#)

[2.2.3. Transient Climate Response to Cumulative Emissions](#)

[2.3. Warming on different emissions scenarios](#)

[2.3.1. Current policy](#)

[2.3.2. All fossil fuels without underground coal gasification](#)

[2.3.3. The worst-case scenario](#)

[2.4. Structural uncertainty](#)

[3. Our future in paleoclimatic context](#)

[3.1. Our climate](#)

[3.2. Hothouses and rapid warming](#)

[3.3. Were past greenhouses inhospitable?](#)

[3.3.1. Pre-Cretaceous warming and mass extinctions](#)

[Warming at the Permian-Triassic boundary](#)

[Volcanic kill mechanisms](#)

[3.3.2. Biodiversity after Pangea](#)

[Cretaceous warming](#)

[The Paleocene-Eocene Thermal Maximum and Eocene Climatic Optimum](#)

[Miocene Climatic Optimum](#)

[Mid-Pliocene Warm Period](#)

[The Last Interglacial](#)

[Transition into the Holocene from last glacial](#)

[The Quaternary Megafaunal Extinction](#)

[3.4. Hominid flourishing and climate change](#)

[3.5. Why did things change after Pangea?](#)

[3.5.1. Pangea](#)

[3.5.2. What about prior to Pangea?](#)

[3.6. Summarising lessons from the paleoclimate](#)

[4. Economic and demographic trends](#)

[4.1. Economic trends](#)

[4.1.1. Global growth](#)

[4.1.2. Regional growth](#)

[Regional growth and discounting](#)

[4.1.3. Growth, impacts, adaptation and course correction](#)

[Subsequent collapse](#)

[4.2. Demographic trends](#)

[5. Agriculture](#)

[5.1. Context and trends in agriculture](#)

[5.1.1. Food production](#)

[5.1.2. Agriculture and the economy](#)

[5.1.3. Historical trends in water use and irrigation](#)

[Water use](#)

[Irrigation](#)

[5.2. By what mechanism does climate change affect food production?](#)

[5.3. What is the projected effect of climate change on agricultural production?](#)

[5.3.1. Crop yields](#)

[Correlated yield declines](#)

[5.3.2. Crop migration](#)

[5.3.3. Water stress and agriculture](#)

[The effect of climate change on precipitation](#)

[The effect of climate change on drought](#)

[What effect will climate-induced water stress have on agriculture and the economy?](#)

[Irrigation is an effective adaptation measure](#)

[Water management is poor](#)

[Desalination](#)

[What are the economic costs of water stress?](#)

[5.3.4. Heat stress for agricultural workers](#)

[5.4. Future agricultural progress](#)

[5.5. Overall judgement on the effects of climate change on agriculture](#)

[5.6. Extreme warming and agriculture](#)

[5.6.1. Lethal limits](#)

[At what level of warming would plants pass lethal limits?](#)

[Europe](#)

[North America](#)

[China](#)

[Russia](#)

[India](#)

[5.6.2. Extreme heat stress and agricultural labour](#)

[5.6.3. Summary of extreme warming and agriculture](#)

[5.7. Ecosystem collapse and threats to agriculture](#)

[5.7.1. Trends in biodiversity](#)

[Species extinctions](#)

[Population losses](#)

[Speciation and hybridisation](#)

[5.7.2. Drivers of species extinctions](#)

[The Species-Area Relationship](#)

[Climate change as a driver](#)

[5.7.3. Could ecosystem collapse destroy agriculture?](#)

[From human intervention to the destruction of agricultural ecosystem services](#)

[Extinctions have mainly occurred on islands](#)

[Forest cover is increasing in many regions](#)

[Local species richness is stable](#)

[Track record of predictions based on species area relationships](#)

[Models vs the paleoclimate](#)

[Global species loss and agriculture](#)

[Clarifying the argument](#)

[Closeness to pre-modern ecology and living standards are negatively correlated](#)

[Overall judgement on ecosystem collapse and agriculture](#)

[6. Heat stress](#)

[6.1. Metrics](#)

[6.2. Heat stress today](#)

[6.3. What are high levels of heat stress actually like?](#)

[6.4. Impacts](#)

[6.4.1. The scope for adaptation](#)

[6.4.2. Death estimates](#)

[6.5. The future human climate niche](#)

[7. Sea level rise](#)

[7.1. Past trends in relative sea level rise and flooding](#)

[7.1.2. Vertical land movement](#)

[7.1.3. Floods](#)

[7.2. How much will sea level rise due to global warming?](#)

[7.3. What is the scope to adapt to relative sea level rise?](#)

[7.3.1. Adaptation options](#)

[7.3.2. Future adaptation prospects](#)

[Up to 2 metres over 80 years](#)

[Relative sea level rise of 5 metres over 100 years](#)

[Past experience with subsidence and living below sea level](#)

[Future projections of adaptation](#)

[7.3.3. Effects on 21st Century sea level rise on land loss and migration](#)

[7.3.4. Economic effects of sea level rise](#)

[7.4. Overall verdict on sea level rise](#)

[8. Tipping points](#)

[8.1. Permafrost carbon release](#)

[8.2. Methane clathrates](#)

[8.2.1. Methane emissions so far](#)

[8.2.2. Most methane would not reach the atmosphere](#)

[8.2.3. The paleoclimate record](#)

[8.2.4. Expert estimates of emissions from methane hydrates](#)

- [8.3. Amazon forest dieback](#)
- [8.4. Boreal forest dieback](#)
- [8.5. Collapse of the Atlantic Meridional Overturning Circulation](#)
 - [8.5.1. Climatic effects of AMOC collapse](#)
 - [8.5.2. How would AMOC collapse affect human society?](#)
 - [8.5.3. The probability of AMOC collapse](#)
- [8.6. Changes to the West African monsoon](#)
- [8.7. Indian monsoon shift](#)
- [8.8. Cloud feedbacks](#)
 - [8.8.1. How plausible are rapid cloud feedbacks?](#)
 - [8.8.2. What effects would rapid cloud feedbacks have on human society?](#)
- [8.9. Hothouse Earth](#)
- [9. Greenhouses and habitability](#)
 - [9.1. The long-term habitability of Earth](#)
 - [9.2. Could humans cause a moist greenhouse or runaway greenhouse?](#)
 - [9.3. How might a moist or runaway greenhouse kill us?](#)
 - [9.4. Quantifying the direct extinction risk of climate change](#)
- [10. Economic costs](#)
 - [10.1. Context and trends](#)
 - [10.2. The social cost of carbon](#)
 - [10.3. Estimating climate damages](#)
 - [10.4. Bottom-up studies](#)
 - [10.4.1. Microeconomic evidence of how temperature affects productivity](#)
 - [10.4.2. Included sectors and impacts](#)
 - [10.4.3. Tipping points](#)
 - [AMOC collapse](#)
 - [West African Monsoons](#)
 - [Indian monsoon shift](#)
 - [Collapse of the West Antarctic Ice Sheet](#)
 - [Rapid cloud feedbacks](#)
 - [How do tipping points affect the economic costs of climate change?](#)
 - [10.4.4. Indirect effects](#)
 - [10.4.5. Regional or global aggregation](#)
 - [10.4.6. Growth or levels?](#)
 - [10.4.7. Overall judgement on bottom-up studies](#)
 - [10.5. Top-down studies](#)
 - [10.5.1. Cross-sectional correlations](#)
 - [Cross-country correlations](#)
 - [The cross-country correlation between temperature and per capita GDP](#)
 - [The cross-country correlation between temperature and growth rates](#)
 - [Methodological issues with cross-country regressions](#)
 - [Cross-sectional relationship within countries](#)
 - [Methodological issues with within-country cross-sectional regressions](#)

- [Other problems with cross-sectional evidence](#)
 - [Overall judgement on the cross-sectional evidence](#)
 - [10.5.2. Weather and economic performance](#)
 - [Methodological drawbacks of weather studies](#)
 - [Weather is different to climate change](#)
 - [Noisy proxies](#)
 - [Non-market impacts](#)
 - [Findings of weather studies](#)
 - [GDP growth or GDP levels?](#)
 - [Adaptation](#)
 - [Model uncertainty](#)
 - [Overall judgement on weather studies](#)
 - [10.6. Expert elicitation](#)
 - [10.6.1. Methodological problems with expert elicitation](#)
 - [10.6.2. Findings of expert elicitation studies](#)
 - [10.6.3. Overall judgement on expert elicitation studies](#)
 - [10.7. Overall view on climate damages](#)
 - [10.8. Climate economics and the long-term](#)
 - [10.8.1. Risk of direct extinction or collapse](#)
 - [10.8.2. Size of indirect risks](#)
 - [10.8.3. Long-run stagnation](#)
 - [10.8.4. The relative geopolitical power of different countries](#)
 - [10.8.5. What is the sign of the effect of economic growth on the long-term?](#)
- [11. Displacement and migration](#)
 - [11.1. Definitions and trends](#)
 - [11.1.1. Displacement](#)
 - [11.1.2. Migration](#)
 - [International migrants](#)
 - [Internal migrants](#)
 - [Migration to coastal regions](#)
 - [11.1.3. The character of migration and displacement](#)
 - [11.2. How will climate change affect displacement?](#)
 - [11.2.1. Coastal displacements from storms and floods](#)
 - [11.2.2. River flood displacements](#)
 - [11.2.3. Overall displacements from floods and storms](#)
 - [11.3. How will climate change affect migration?](#)
 - [Myers estimates](#)
 - [Other estimates](#)
 - [Summary of migration estimates](#)
 - [11.4. Summary of climate change, migration and displacement](#)
- [12. Conflict](#)
 - [12.1. Context and trends](#)
 - [12.1.1. Definitions](#)

- [12.1.2. Trends](#)
- [12.1.3. Drivers of the Long Peace](#)
- [12.2. Climate change and pre-industrial strife](#)
- [12.3. The contemporary relationship between climate change and conflict](#)
 - [12.3.1. The correlation](#)
 - [12.3.2. Overview of the literature](#)
 - [12.3.4. Effect sizes found in the literature](#)
 - [12.3.5. A review of Burke et al \(2009\)](#)
 - [Failure to predict decline in civil conflict incidence](#)
 - [Sensitivity to civil war cut-off](#)
 - [Aggregated spatial and temporal scales](#)
 - [Qualitative analysis](#)
 - [Overall view](#)
 - [12.3.6. A review of the Burke et al \(2015b\) meta-analysis](#)
 - [Intergroup conflict includes a lot](#)
 - [Study inclusion criteria](#)
 - [Study inclusion and coding decisions](#)
 - [Reporting bias](#)
 - [Implausibly precise estimates](#)
 - [Africa focus](#)
 - [Caruso et al \(2016\) - Climate Change, Weather Shocks and Violent Conflict: A Critical Look at the Evidence](#)
 - [Dell et al \(2012\) - Temperature Shocks and Economic Growth: Evidence from the Last Half Century', American Economic Journal: Macroeconomics](#)
 - [Hsiang et al \(2011\) - Civil conflicts are associated with the global climate, Nature](#)
 - [Other issues with panel studies](#)
 - [Overall judgement on these estimates](#)
 - [Implications of the estimates](#)
 - [Burke et al estimates taken at face value](#)
 - [My favoured estimates](#)
 - [Mechanisms of impact](#)
 - [12.3.7. Climate less important than other factors](#)
 - [Future conflict predictions](#)
 - [Hegre et al \(2016\) 'Forecasting Civil Conflict along the Shared Socioeconomic Pathways', Environmental Research Letters](#)
 - [Witmer et al \(2017\) 'Subnational violent conflict forecasts for sub-Saharan Africa, 2015–65, using climate-sensitive models', Journal of Peace Research](#)
- [12.4. Concluding thoughts on climate change and civil conflict](#)
- [12.5. Climate, conflict and Great Power War](#)
 - [12.5.1. How would Great Power War affect the long-term future?](#)
 - [12.5.2. Theories of interstate war](#)
 - [Realism](#)
 - [Rational states?](#)
 - [Liberalism](#)

[Constructivism](#)

[Public choice theory](#)

[12.5.3. How might climate change affect the risk of Great Power War?](#)

[Conflict over water resources](#)

[Economic costs](#)

[Civil conflict and upheaval](#)

[Mass migration, displacement and conflict](#)

[The mechanism](#)

[The size of the effect on conflict](#)

[12.5.4. The main drivers of interstate war](#)

[War usually involves contiguous neighbours](#)

[A large proportion of wars involve enduring or strategic rivals](#)

[Security dilemma and conflict spiral](#)

[Unresolved territorial disputes](#)

[Misperceptions about relative power](#)

[Power transitions](#)

[12.5.4. The influence of climate change on the risk of Great Power War](#)

[The US and China](#)

[The US and Russia](#)

[India and Pakistan](#)

[Water wars](#)

[Economic performance](#)

[Mass migration and civil conflict](#)

[India and China](#)

[12.5.5. Which conflicts would do the most damage?](#)

[AI](#)

[Biorisk](#)

[Nuclear weapons](#)

[Other technological risks](#)

[12.5.6. Quantifying the indirect risks of climate change](#)

[12.5.7. Overall judgement on climate change and Great Power War](#)

[13. Interpersonal conflict, aggression and violence](#)

[13.1. Context, base rates and trends](#)

[13.2. Cross-sectional relationship](#)

[13.3. Experimental and quasi-experimental evidence](#)

[13.3.1. Methodological issues](#)

[Weather different to climate](#)

[Heavy US focus](#)

[Implausible precision](#)

[Reporting bias](#)

[13.4. Mechanisms](#)

[13.5. Effect size](#)

[Appendix 1. The WCRP estimate of climate sensitivity](#)

[Prior selection](#)

[Using multiple lines of evidence](#)

[Formally producing a posterior distribution](#)

Introduction

In this report, I will evaluate the scale of climate change from a *longtermist* point of view. [Longtermism](#) is the idea that influencing the long-term future, thousands of years into the future and beyond, is a key moral priority of our time.

The basic case for longtermism is as follows. First, future sentient life matters. Our lives surely matter just as much as those lived thousands of years ago — so why shouldn't the lives of people living thousands of years hence matter equally? Second, the future could be vast. Absent catastrophe, most people who will ever live have not yet been born. Third, our actions may predictably influence how well this long-term future goes. In sum, we have a responsibility to ensure that future generations get to survive and flourish.

Climate change is a proof of concept of longtermism. Every time we drive, fly, or flick a light switch, each of us causes CO₂ to be released into the atmosphere. This changes the amount of CO₂ that is in the atmosphere for a very long time: unless we suck the CO₂ out of the atmosphere ourselves, concentrations only fall back to natural levels after hundreds of thousands of years. The chart below shows long-term CO₂ concentrations after different amounts of cumulative carbon emissions.

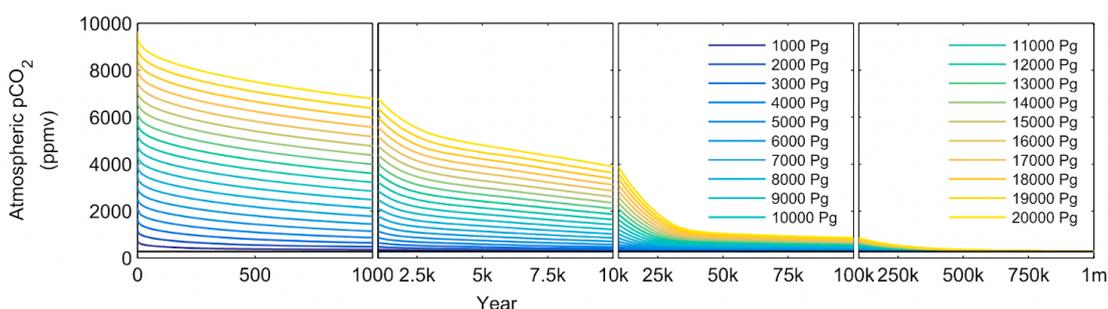


Figure 2. Atmospheric pCO₂ predicted by cGENIE for the pulse series scenarios (1000–20,000 Pg C). Preindustrial CO₂ concentrations are shown in black.

Source: N. S. Lord et al., 'An Impulse Response Function for the "Long Tail" of Excess Atmospheric CO₂ in an Earth System Model', *Global Biogeochemical Cycles* 30, no. 1 (2016): 2–17, <https://doi.org/10.1002/2014GB005074>.

Because CO₂ is a greenhouse gas, it traps heat leaving the Earth's surface, causing the planet to warm.

Some of the ecological effects of climate change get worse over time. The clearest example of this is sea level rise. On current policy, the most likely sea level this century is 75cm. However, over 10,000 years, sea levels will rise by 10 metres. Over the long-term, the world will look very different.

Climate change is a long-term problem in which we are all implicated; by emitting CO₂, we contribute to a problem that imposes costs on our descendants for generations to come.

Some economists think it is appropriate to discount climate impacts that occur in the future. They do this for two reasons. Firstly, they assume that future generations will be richer and so will be better able to deal with climate impacts. It is not clear that this is justified. Climate impacts will be especially bad for developing countries at low latitude, which will account for a large fraction of the human population by the end of the century. Historically, many countries, especially African ones, have had low rates of economic growth, so it is not clear that they will be much richer in the future. Even though average living standards will improve, these improvements are likely to be concentrated in countries outside of Africa, and a substantial fraction of people may not be much better off.

Secondly, some economists defend positive *pure time preference*, which is the view that we should value the welfare that occurs in the future less because it is in the future. For example, Nordhaus (2008) uses a 1% rate of pure time preference.³ This implies that one death today is worth more than 100 deaths in 500 years. Economists defend pure time preference on the basis that, as shown in people's saving and consumption behaviour, people in the real world are impatient and prefer a benefit today to one in a year's time. It is difficult to see why people's impatience about *their own* welfare is relevant to how much weight to put on the welfare of other people who will live in the future. It is, for instance, exactly analogous to arguing that my child's future welfare is worth less than my own on the sole basis that I have frittered away their inheritance at a theme park.

Disagreements about the weight to put on the welfare of future generations can have crucial importance for policy. The more weight you put on the welfare of future generations, the more likely you are to favour aggressive climate policy. Indeed, disagreement about the discount rate largely explains why the UK government's Stern Review favoured much more aggressive climate policy than Nordhaus' prominent DICE model.

From a longtermist point of view, it is especially important to avoid outcomes that could have *persistent* and *significant* effects. These include events like human extinction, societal collapse, a permanent negative change in human values, or prolonged economic stagnation. If we go extinct, then that would be the end of the human story, and there would be no future generations at all. If civilisation collapses permanently, then future generations will be left much worse off than they could have been, living lives full of suffering rather than ones of flourishing.

³ Nordhaus, W. (2008). *A question of balance*. New Haven, CT: Yale University Press.

1. How much will we emit?

As Stefan Schubert argues in his piece on [Sleepwalk bias](#), there are two different kinds of prediction:

1. What I bet will happen
2. Predictions as warnings - if we don't do x (get our act together), then y will happen.

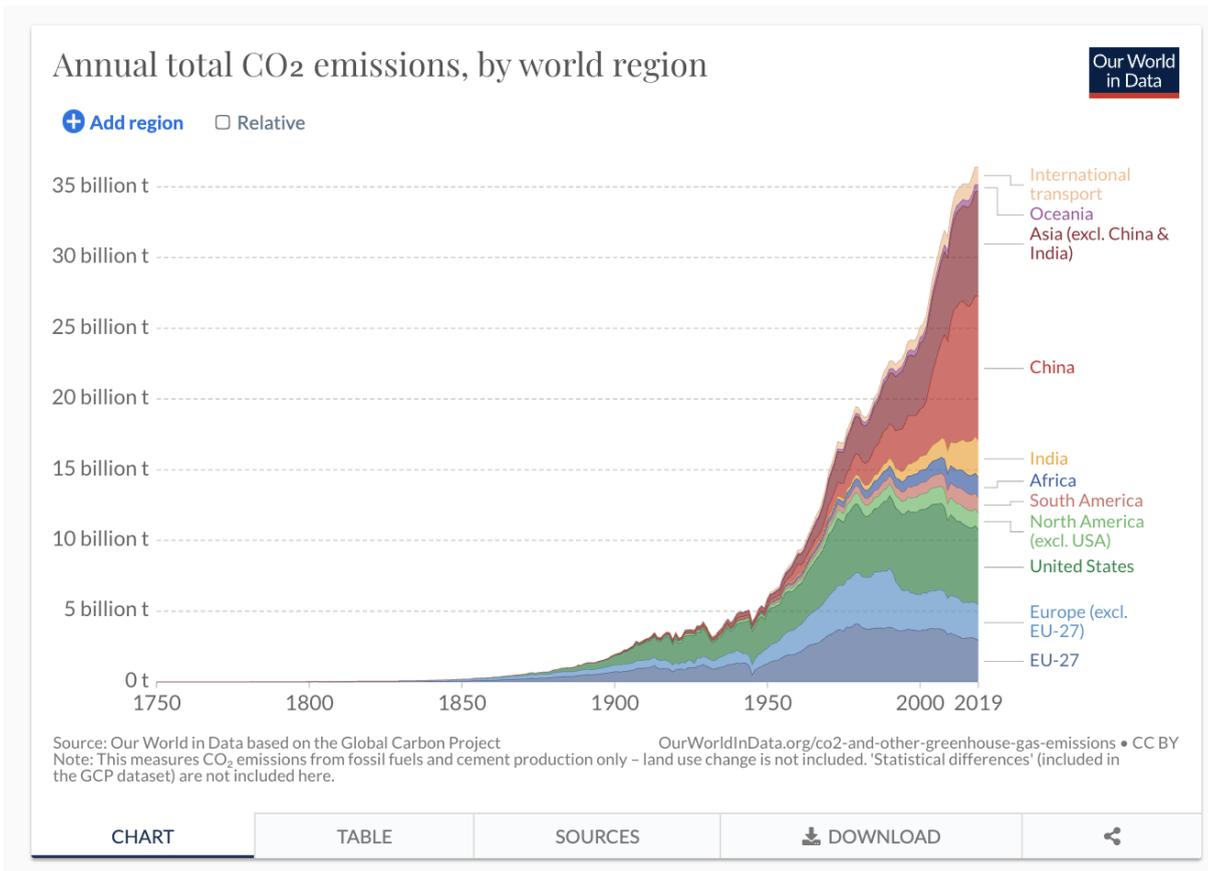
There are few probabilistic assessments of the emissions pathways (type 1 above) and a lot of the scenarios are more like warnings, which are conditional probabilities of different kinds, such as:

1. If current policy never strengthens, this is where we could go
2. If current pledges and promises are followed, this is where we could go
3. If trends continue as they have been doing, this is where emissions could go.

From the perspective of risk management, we are more interested in the bets than the warnings. We want to know how likely different emissions scenarios are, not what emissions scenarios are most likely conditional on an improbable set of assumptions about technology and policy.

1.1. Emissions so far

James Watt's patent for the steam engine in 1769 marked the start of the Industrial Revolution and centuries of almost unchecked fossil fuel burning. Prior to COVID, we were putting about 36 billion tonnes of CO₂ into the atmosphere every year.



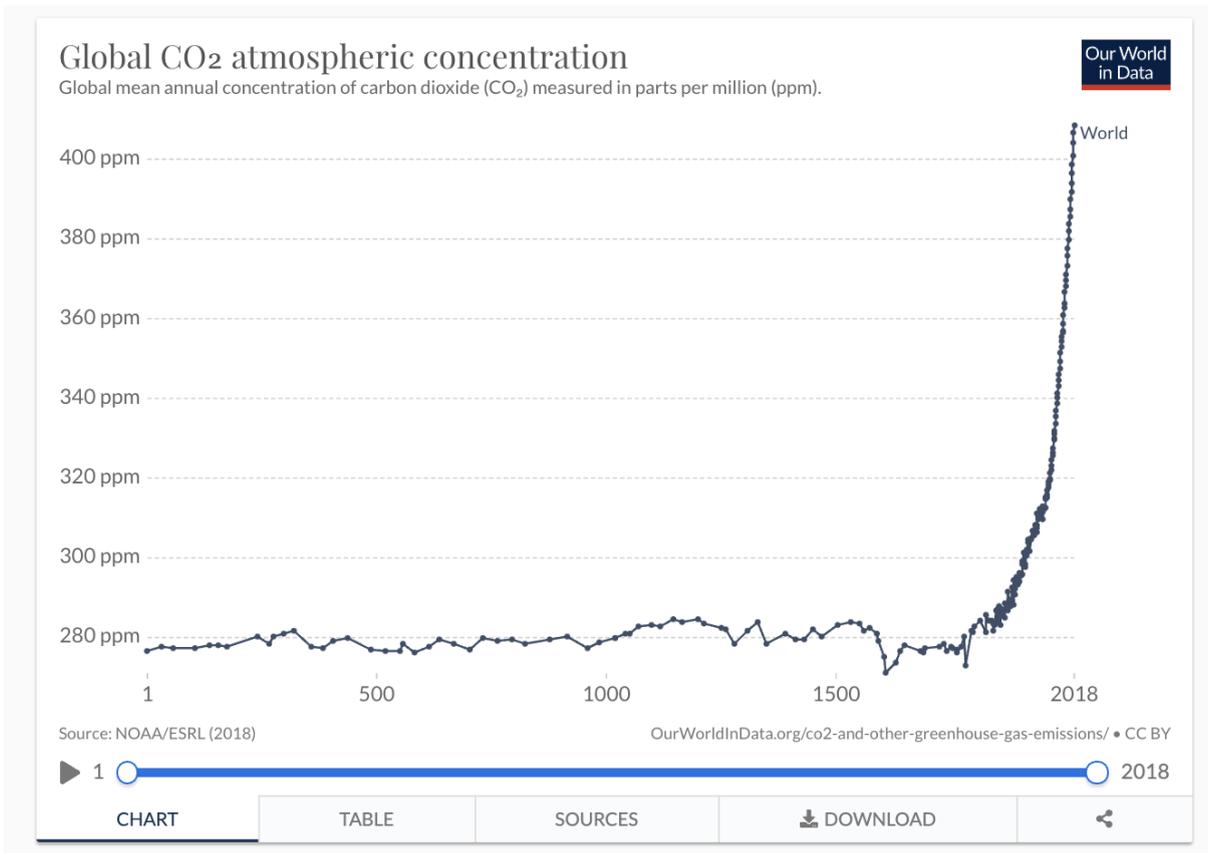
Despite increasing attention to climate change, more than half of all CO₂ ever emitted was released [after 1990](#). Since the 1980s, much of the growth in emissions has been driven by emerging economies in Asia. In the future, most energy demand and emissions growth is likely to come from [outside Europe and North America](#).

Between 1750 and 2017, [cumulative emissions from fossil fuels and cement](#) were 1,580 billion tonnes of CO₂ (gigatonnes of CO₂ (GtCO₂)). This is 431 GtC (because 1 GtC = 3.667 GtCO₂).⁴ Note that a gigatonne is the same as a petagram (Pg), another metric sometimes used in climate science.

Unless we eventually remove CO₂ from the atmosphere ourselves, CO₂ accumulates in the atmosphere for millennia. This means that as long as CO₂ emissions are positive, concentrations of CO₂ in the atmosphere will continue to rise.

Immediately prior to the Industrial Revolution, CO₂ concentrations were 278 parts per million. They are now at 415ppm

⁴ Note that this does not include cumulative emissions from land use change.



At current emissions rates, CO₂ concentrations are rising at about [2.5 ppm each year](#).

1.2. Future emissions

The latest IPCC report outlines a range of emissions scenarios known as 'Representative Concentration Pathways' (RCPs), as well as a range of 'Shared Socioeconomic Pathways' - socioeconomic narratives about how the world will develop in key areas such as population, income, inequality and education. There are 5 Shared Socioeconomic Pathways.

- SSP1: Sustainability (Taking the Green Road)
- SSP2: Middle of the Road
- SSP3: Regional Rivalry (A Rocky Road)
- SSP4: Inequality (A Road divided)
- SSP5: Fossil-fueled Development (Taking the Highway)

These diagrams show some key assumptions of the shared socioeconomic pathways.

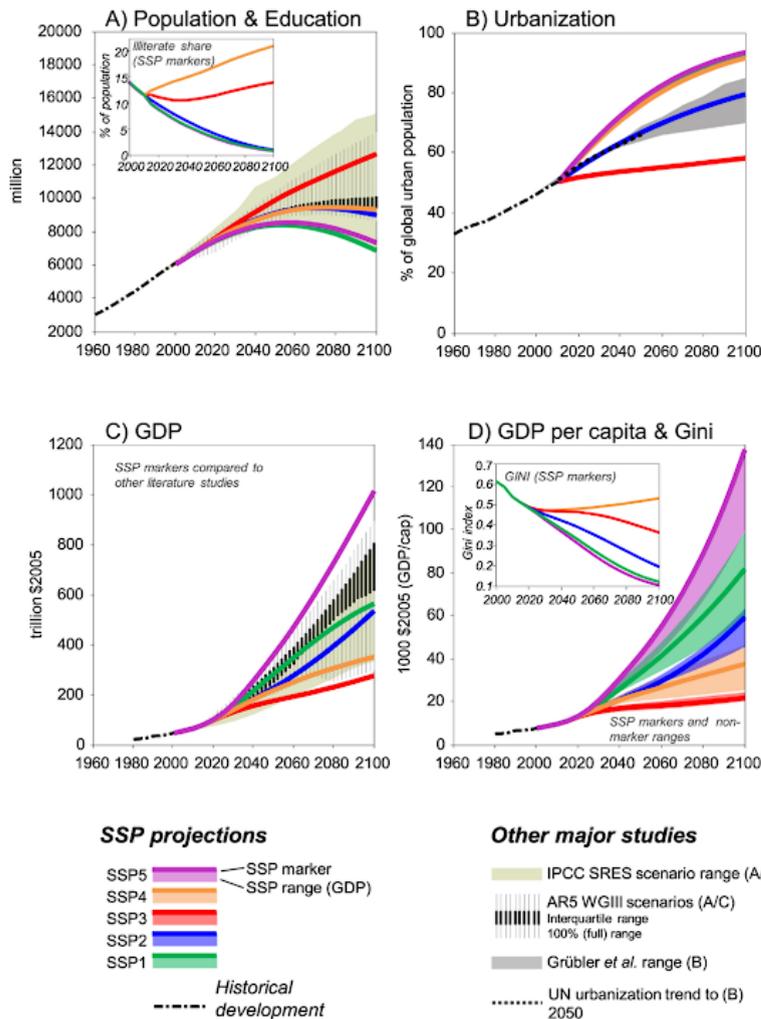


Fig. 2. Development of global population and education (A), urbanization (B), GDP (C), and GDP per capita and the Gini index (D). The inset in panel A gives the share of people without education at age of ≥ 15 years, and the inset in panel D denotes the development of the global (cross-national) Gini index. The SSPs are compared to ranges from other major studies in the literature, such as the IPCC AR5 (Clarke et al., 2014); IPCC SRES (Nakicenovic and Swart, 2000), UN, and Grübler et al. (2007). The colored areas for GDP (panel D) denote the range of alternative SSP GDP projections presented in this Special Issue (Dellink et al. (2016), Crespo Cuaresma (2016), Leimbach et al. (2016)).

Source: Keywan Riahi et al., 'The Shared Socioeconomic Pathways and Their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview', *Global Environmental Change* 42 (1 January 2017): 153–68, <https://doi.org/10.1016/j.gloenvcha.2016.05.009>.

SSP2 is meant to be the 'current trends continue' scenario. SSP5 is a scenario involving rapid growth in energy demand and in GDP. You can read the narratives in full [here](#). These scenarios are not all meant to have equal probability so we shouldn't see the spread across SSPs as a probability range.

The chart below shows regional GDP growth on the different SSPs:

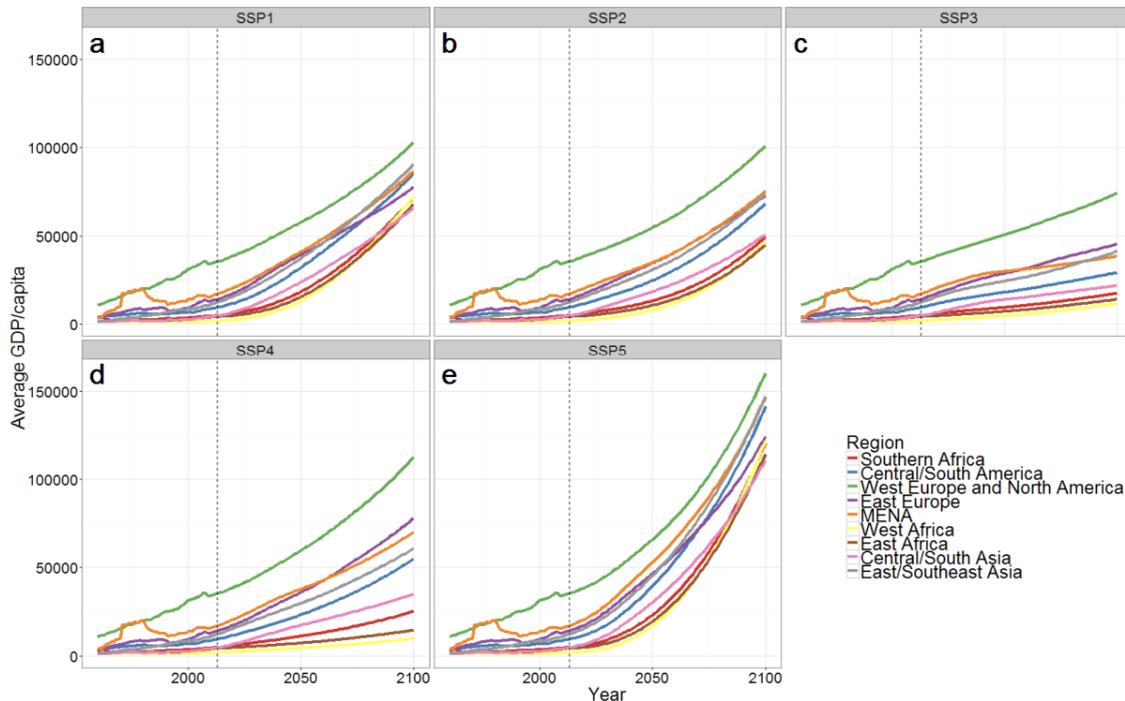


Figure S3. Country average GDP per capita (2005 USD PPP) by region and SSP.

Source: Håvard Hegre et al., 'Forecasting Civil Conflict along the Shared Socioeconomic Pathways', *Environmental Research Letters* 11, no. 5 (April 2016): <https://doi.org/10.1088/1748-9326/11/5/054002>.

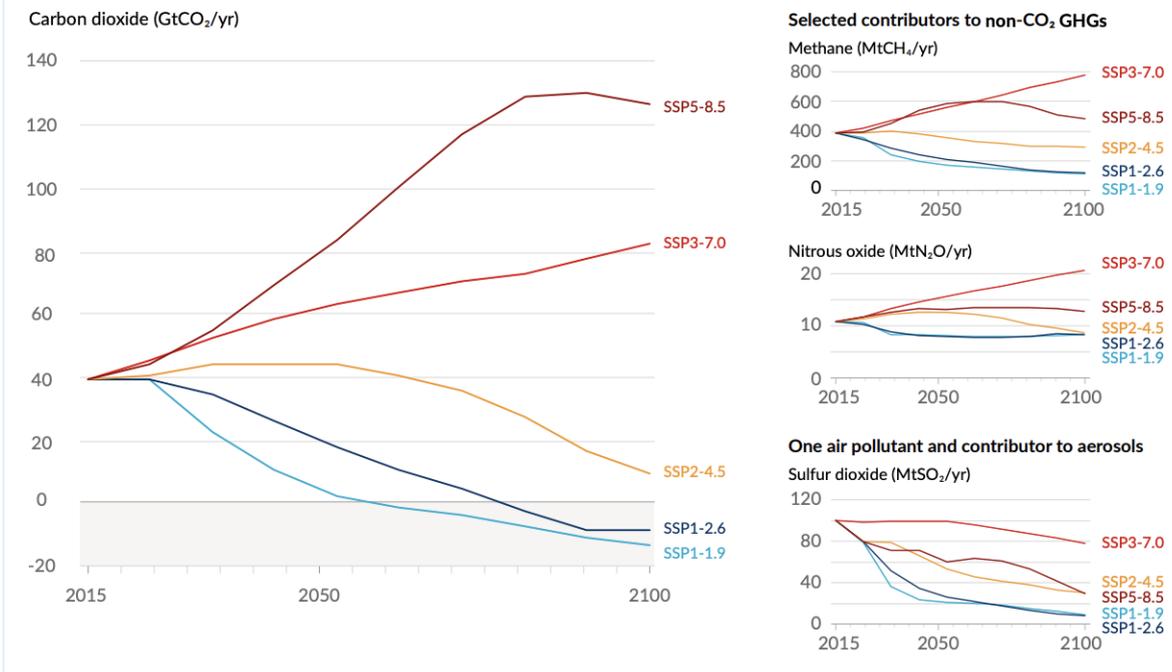
As I discuss in chapter 5, given historical trends, the most plausible SSP seems to be SSP3 or SSP4. The other SSPs predict much higher growth in Africa than is warranted by historical experience so far.

Each Shared Socioeconomic Pathway can be combined with a range of different Representative Concentration Pathways. For instance, the high growth Shared Socioeconomic Pathway 5 (SSP5) can be combined with modest and also very high emissions, depending on assumptions about the strength of climate policy.

The chart below shows a subset of shared socioeconomic pathways. The scenario labels combine the SSP with the Representative Concentration Pathway. So, for instance, SSP5 combined with RCP8.5, is called 'SSP5-8.5', while SSP2 combined with RCP4.5 is called 'SSP2-4.5'. The number of each RCP refers to the radiative forcing in Watts per square metre at 2100 on the different emissions scenarios: the higher the number, the greater the warming.

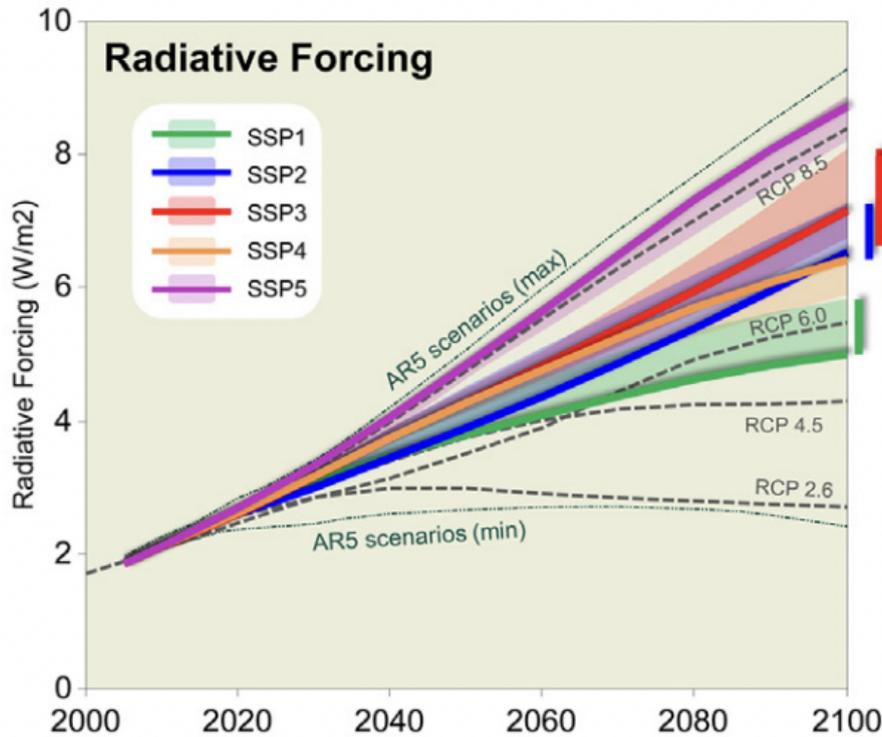
Future emissions cause future additional warming, with total warming dominated by past and future CO₂ emissions

a) Future annual emissions of CO₂ (left) and of a subset of key non-CO₂ drivers (right), across five illustrative scenarios



IPCC, *Climate Change 2021: The Physical Science Basis*, Assessment Review 6, [Summary for Policymakers: Figure SPM.4](#)

There are also ‘baseline’ SSPs which provide a “description of future developments in absence of new climate policies beyond those in place today”, i.e. on current policy. (riahi 155)



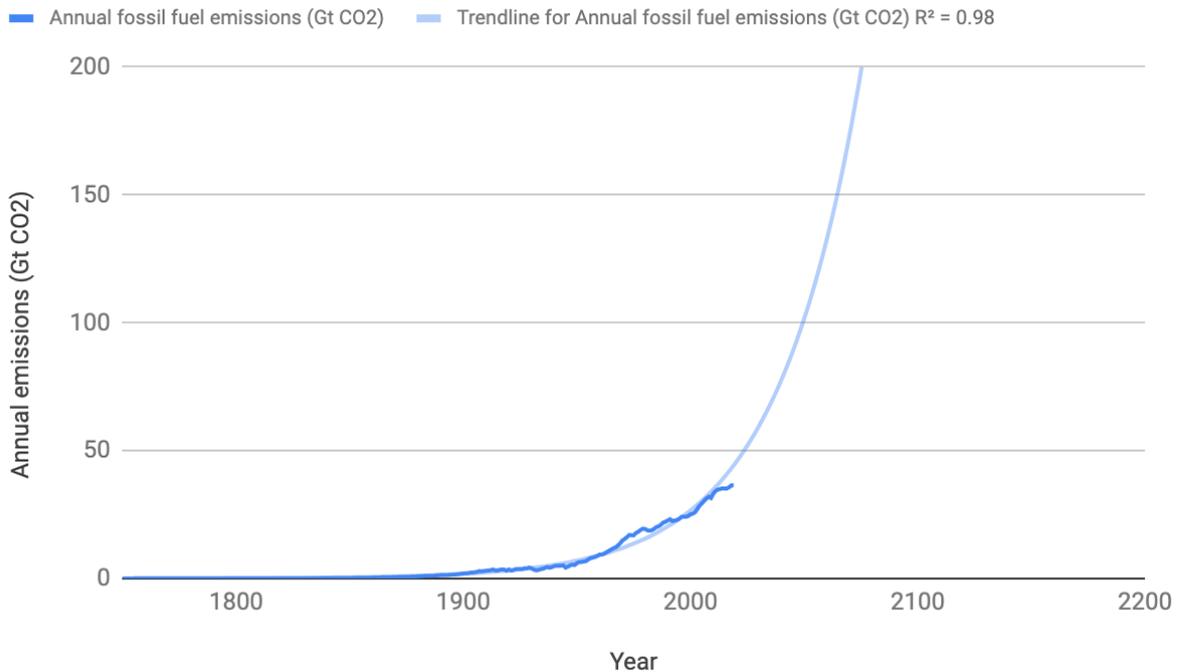
Source: Keywan Riahi et al., 'The Shared Socioeconomic Pathways and Their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview', *Global Environmental Change* 42 (1 January 2017): Fig 5, <https://doi.org/10.1016/j.gloenvcha.2016.05.009>.

For SSP2, 3 and 4 the baseline scenario is around RCP6 to RCP7. For SSP1, it is below RCP6, while for SSP5, it is RCP8.5

To understand the risk posed by climate change, what we would like to know is the probability of different emissions pathways. The IPCC does not put a probability on future emissions scenarios and probabilistic assessments are thin on the ground in the wider literature. However, some recent studies suggest that there is now more cause for optimism than there ever has been before. While many previously assumed that RCP6 or RCP8.5 was the most likely outcome on current policy, it now looks like RCP4.5 is more likely, and that climate policy is going to strengthen in the future.

1.2.1. Climate progress

Since the Industrial Revolution, emissions from fossil fuel burning and cement increased at upwards of 3% per year, on average, and the historical trend fits an exponential closely. If this trend were to continue, emissions would be extremely high by the end of the century, well in excess of SSP5-8.5:



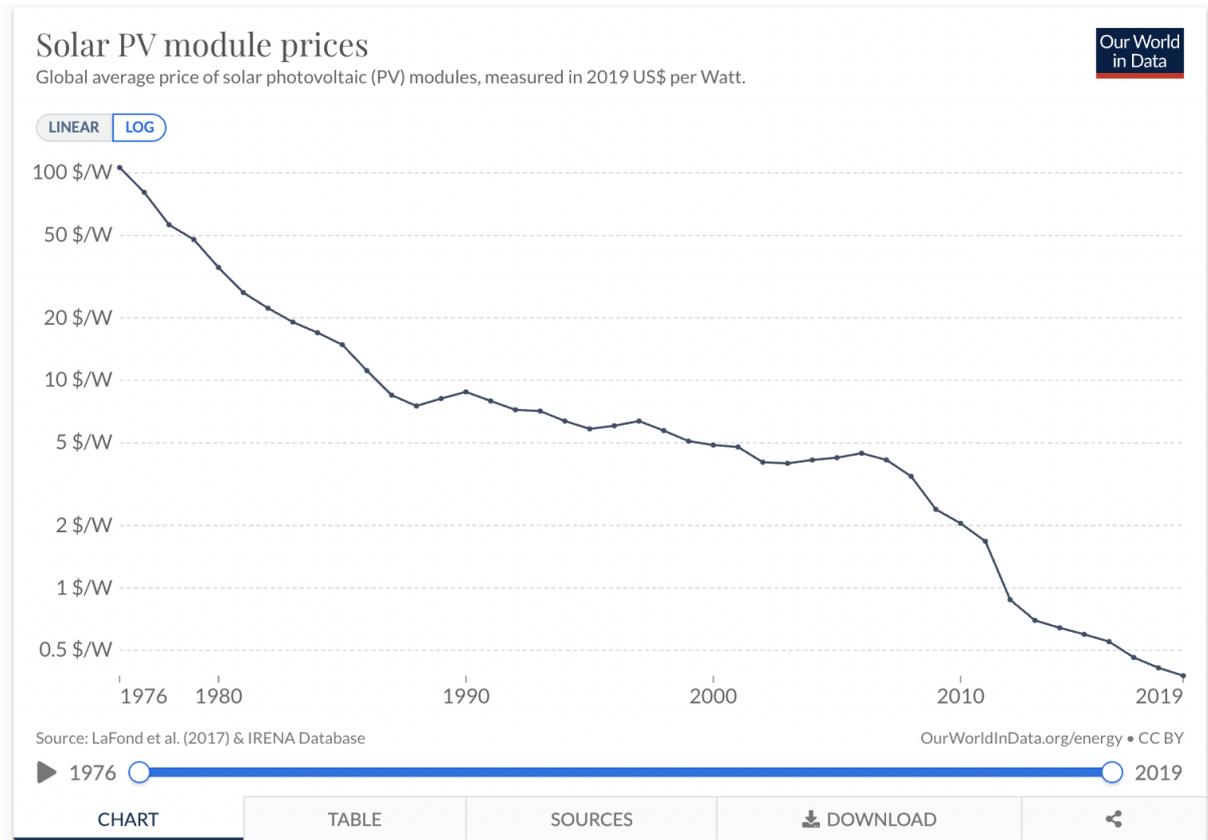
Source: [Our World in Data](#)

Between 2000 and 2010, emissions grew at 2.8% per year. However, the picture on climate change looks more optimistic than it once did. Emissions growth slowed down dramatically between 2010 and 2019, to only around 1%. This decline was (as explained [here](#)) due to improvements in energy efficiency, in the emissions intensity of energy (i.e. switching to less polluting sources of energy), and due to slower than expected economic growth.⁵

Many recent studies suggest that the highest emissions scenarios now look much less likely. There are several reasons for this. Firstly, the costs of renewables and batteries have declined extremely quickly. Historically, models have been too pessimistic on cost declines for solar, wind and batteries: out of nearly 3,000 Integrated Assessment Models, none projected that solar investment costs would decline by more than 6% per year between 2010 and 2020. In fact, they declined by 15% per year.⁶ This chart from Our World in Data shows the decline in the cost of solar module costs, with costs falling by more than 99% over 43 years:

⁵ Matthew G. Burgess et al., 'IPCC Baseline Scenarios Have Over-Projected CO2 Emissions and Economic Growth', *Environmental Research Letters* 16, no. 1 (December 2020): 014016, <https://doi.org/10.1088/1748-9326/abcdd2>

⁶ "Sound energy investments require reliable forecasts. As illustrated in Figure 2(a), past projections of present renewable energy costs by influential energy-economy models have consistently been much too high. ("Projections" are forecasts conditional on scenarios, so we use the terms interchangeably.) The inset of the figure gives a histogram of 2,905 projections by integrated assessment models, which are perhaps the most widely used type of global energy-economy models^{19,20,21,22}, for the annual rate at which solar PV system investment costs would fall between 2010 and 2020. The mean value of these projected cost reductions was 2.6%, and all were less than 6%. In stark contrast, during this period solar PV costs actually fell by 15% per year. Such models have consistently failed to produce results in line with past trends^{3,23}" Way et al, '[Empirically grounded technology forecasts and the energy transition](#)', Oxford Martin School, 2021: p.3.



Fundamentally, existing mainstream economic models of climate change consistently fail to model exponential cost declines, as shown on the chart below. The left pane below shows historical declines in solar costs compared to Integrated Assessment Model projections of costs. The pane on the right shows the cost of solar compared to Integrated Assessment Model assessments of 'floor costs' for solar systems - the lowest that solar could go. Real world solar system costs have consistently smashed through these supposed floors.

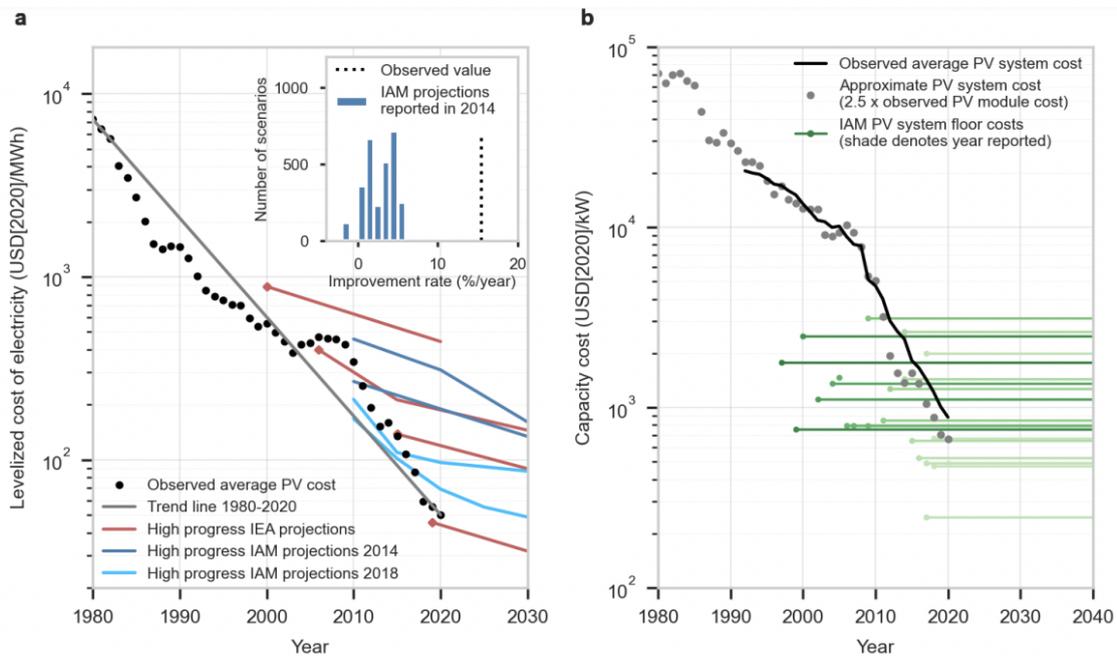


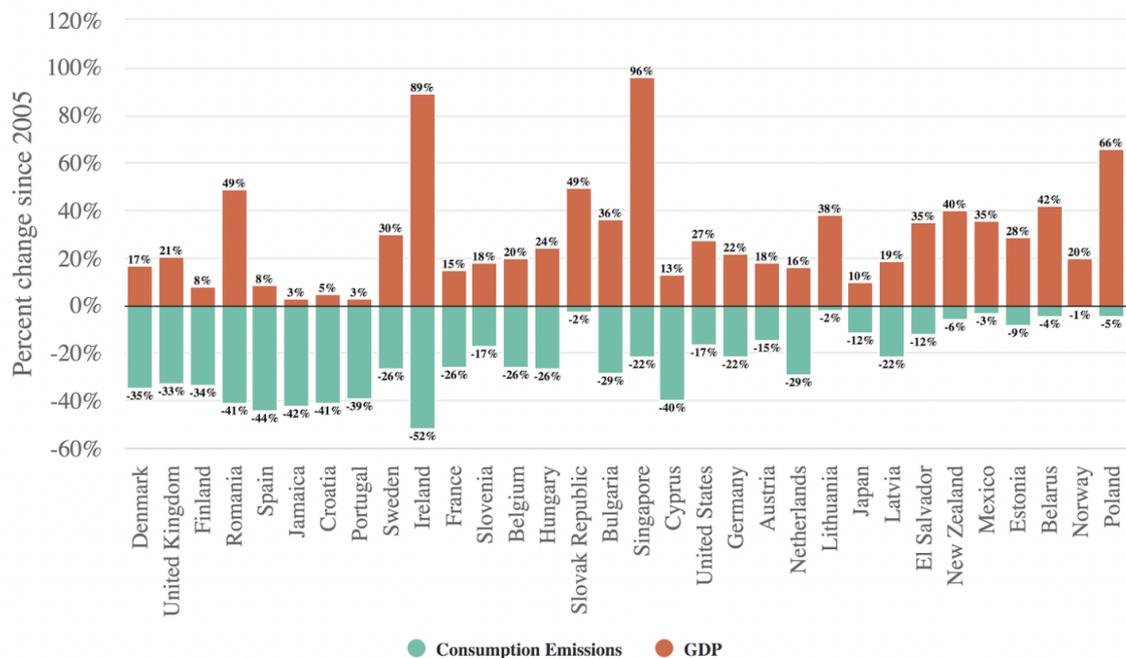
Figure 2: Historical PV cost forecasts and floor costs. (a) The black dots show the observed global average levelized cost of electricity (LCOE) over time. Red lines are LCOE projections reported by the International Energy Agency (IEA), dark blue lines are integrated assessment model (IAM) LCOE projections reported in 2014¹⁹ and light blue lines are IAM projections reported in 2018^{20,21}. IAM projections are rooted in 2010 despite being produced in later years. The projections shown are exclusively “high technological progress” cost trajectories drawn from the most aggressive mitigation scenarios, corresponding to the biggest projected cost reductions used in these models. Other projections made were even more pessimistic about future PV costs. The inset compares a histogram of projected compound annual reduction rates of PV system investment costs from 2010 to 2020 to what actually occurred (based on all 2,905 scenarios for which the data is available¹⁹). (b) PV system floor costs implemented in a wide range of IAMs. The colours denote the year the floor cost was reported, ranging from 1997 (dark green) to 2020 (light green). Observed PV system costs are also shown. The cost of PV modules scaled by a constant factor of 2.5 is provided as a reference. For further details and data sources see Extended Data Figures 6 and 7(a), and SN6.10

Source: Way et al, ‘[Empirically grounded technology forecasts and the energy transition](#)’, Oxford Martin School, 2021.

Secondly, climate policy has strengthened substantially over the last few years. Countries representing [66% of global CO₂ emissions](#) have committed to achieving net-zero emissions by 2050. Most importantly, China has pledged to get to net zero by 2060. Historically, such targets have not been met, but some of them, such as that of the UK, are enshrined in law.

As a consequence of this and other factors such as declining natural gas costs, many countries have recently [absolutely decoupled](#) GDP growth and consumption-based CO₂ emissions. Since the emissions are consumption-based, the decline does not merely reflect offshoring of emissions intensive industries like steel. In a piece for the Breakthrough Institute, Zeke Hausfather shows that over the past 15 years, 32 countries with more than 1 million people have seen rising incomes but lower emissions, even when we adjust for emissions embedded in trade.

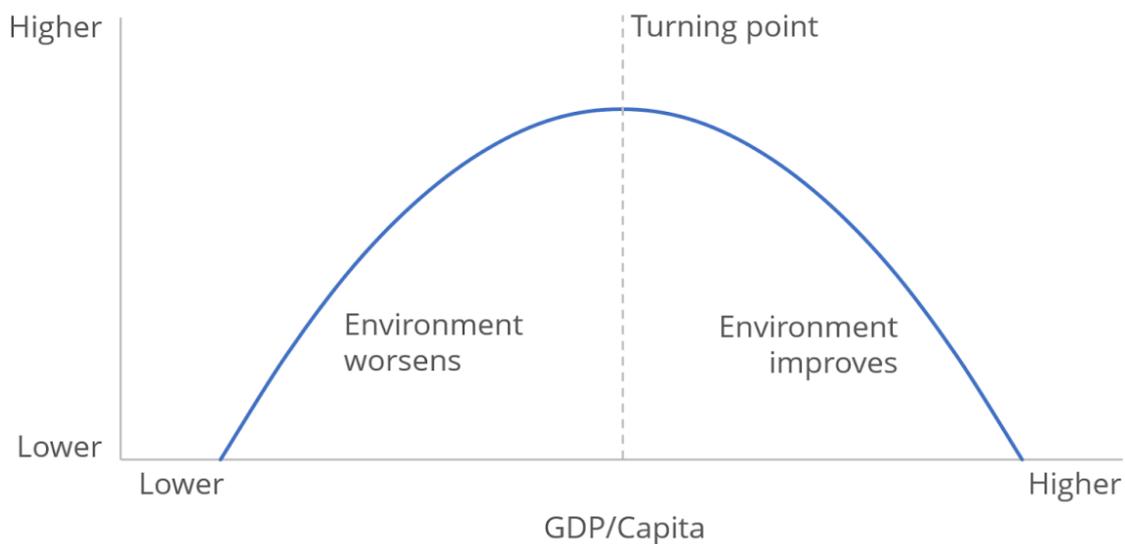
Decoupling of consumption emissions and GDP: 2005-2019



It seems likely that this will be true for more and more countries in the future.

This does not necessarily mean that there is a Kuznets Curve relationship between GDP and CO₂ emissions. On the Kuznets Curve, once income per head increases past a certain point, pollutants decline because environmental protection is a luxury good.

Environmental damage



It is unclear how exactly to define the Kuznets Curve, and identifying one for income and emissions raises some difficult econometric problems. Many studies find there to be an

Environmental Kuznets Curve for carbon, but many don't, and there have been some sharp criticisms of the econometrics used in the studies finding a positive effect.⁷

It is especially difficult to disentangle time effects from income effects in Kuznets Curve studies. Emissions per head are declining over time in many rich countries (which are also experiencing economic growth), but this may be due to the passage of time rather than economic growth. For instance, there might be technological improvements that don't improve growth proportionately, such as less polluting cars that are otherwise identical, or switching from coal to gas, solar or wind, without reducing costs.

To understand what this means - imagine if in Britain, economic growth was expected to be 2% in the coming year and instead it turned out to be 3%. Then we should expect emissions in Britain to be higher than were initially expected. If Britain was past the turning point on the curve, the Environmental Kuznets Curve theory would expect them to be lower. This doesn't contradict that emissions might fall in Britain. On one interpretation of the Kuznets Curve theory they'd fall more if there was more growth and on the opposing theory they would fall less if there was more growth.

As an illustration of this, Italy has had zero income growth for twenty years, but emissions per person have dropped a lot over that time. It would be easy to confuse this time effect with an income effect if Italy had indeed experienced income growth.

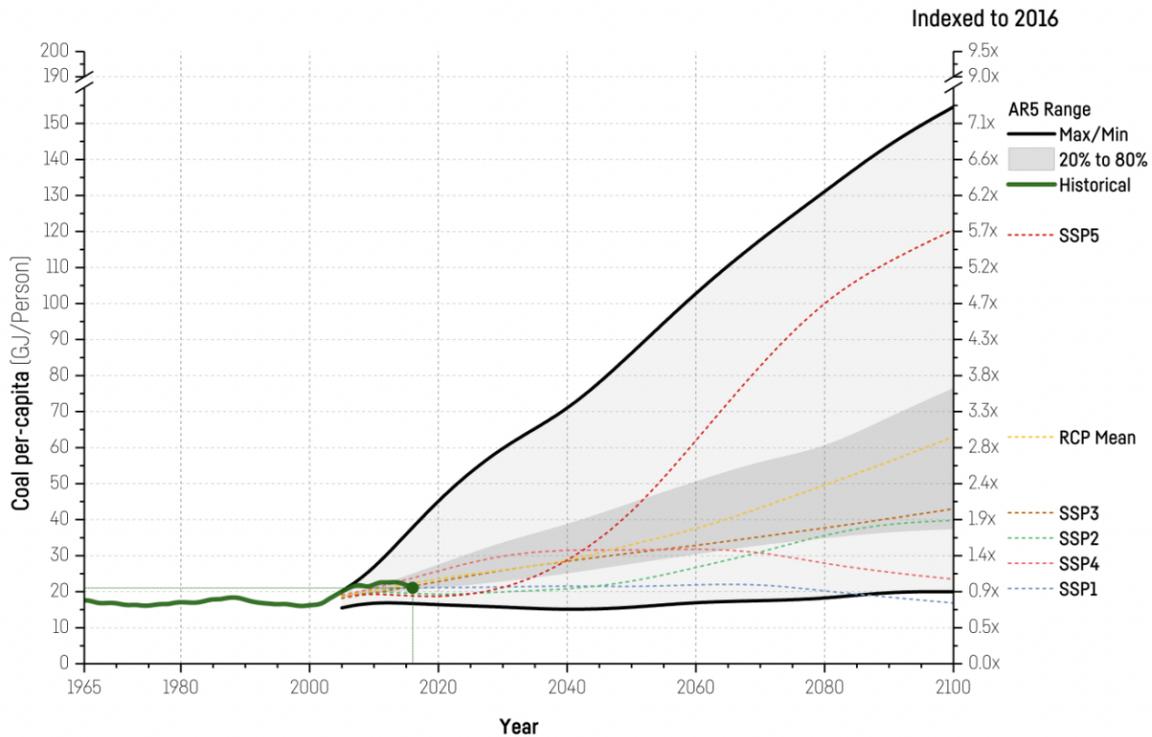
It is unclear how one should specify the correct lag from growth to the Kuznets Curve turning point, i.e. whether we should expect emissions to decline instantaneously, or only a few years later.

For all this, the evidence of absolute decoupling is encouraging.

1.2.2. Assumptions about coal use

To understand SSP 'current policy' models better, it is useful to investigate their assumptions about coal use. Many SSPs assume that there will be a large increase in per capita coal use. This looks very unlikely in part due to the decline in the cost of renewables and the abundance of natural gas driven by hydraulic fracturing, and in part because countries tend to transition away from coal as they get richer. For example, here is the increase in per capita coal use on different baseline SSPs:

⁷ David I. Stern, 'The Environmental Kuznets Curve after 25 Years', *Journal of Bioeconomics* 19, no. 1 (1 April 2017): 7–28, <https://doi.org/10.1007/s10818-017-9243-1>; Herman RJ Vollebergh, Bertrand Melenberg, and Elbert Dijkgraaf, 'Identifying Reduced-Form Relations with Panel Data: The Case of Pollution and Income', *Journal of Environmental Economics and Management* 58, no. 1 (2009): 27–42; Martin Wagner, 'The Environmental Kuznets Curve, Cointegration and Nonlinearity', *Journal of Applied Econometrics* 30, no. 6 (2015): 948–67, <https://doi.org/10.1002/jae.2421>.



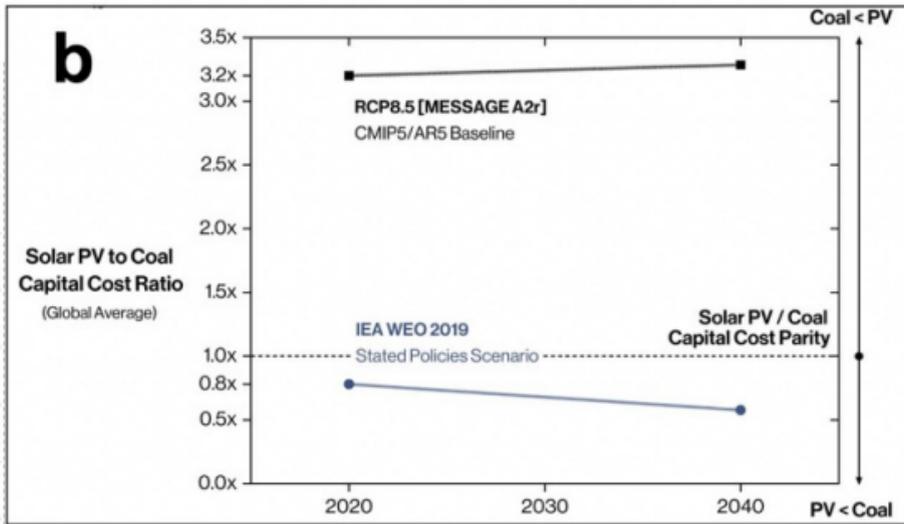
Source: Justin Ritchie and Hadi Dowlatabadi, 'Why Do Climate Change Scenarios Return to Coal?', *Energy* 140 (1 December 2017): fig. 3, <https://doi.org/10.1016/j.energy.2017.08.083>.

For comparison, China burned what was widely seen to be a prodigious amount of coal from 2000 onwards, but that is dwarfed by the increase in coal use projected on SSP5-8.5. Ritchie and Dowlatabadi (2017) argue that the SSPs rely on a coal cost projection model, which assumes that costs decline as we extract more. This has been invalidated by the historical data.⁸

This is illustrated in the figure below from Burgess et al (2020), which shows the assumptions of MESSAGE, a leading integrated assessment model, about solar costs compared to coal costs over the next 30 years. As this shows, MESSAGE assumes that coal capital costs will be threefold cheaper than solar for the next thirty years, even though solar costs are already lower.

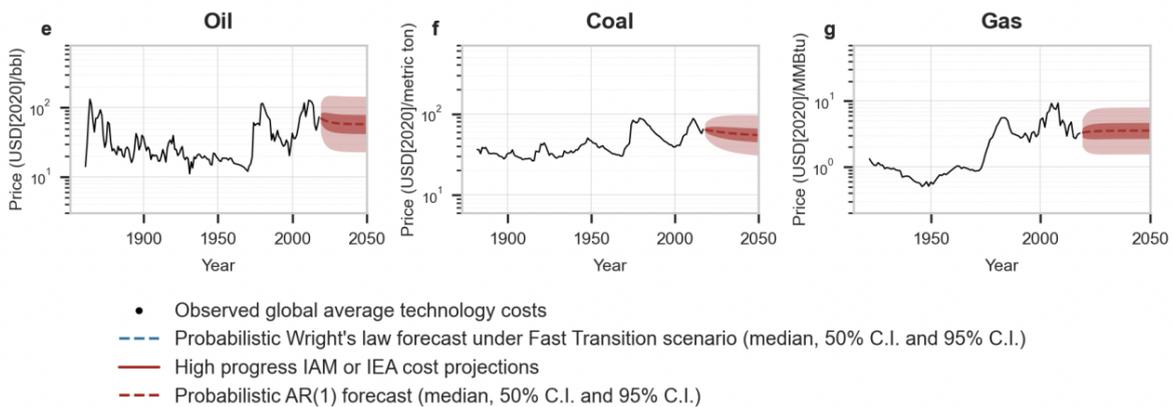
⁸ "To understand possibilities for energy resources in this context, the research community draws from Rogner (1997) which proposes a theory of learning-by-extracting (LBE). The LBE hypothesis conceptualizes total geologic occurrences of oil, gas, and coal with a learning model of productivity that has yet to be empirically assessed.

This paper finds climate change scenarios anticipate a transition toward coal because of systematic errors in fossil production outlooks based on total geologic assessments like the LBE model. Such blind spots have distorted uncertainty ranges for long-run primary energy since the 1970s and continue to influence the levels of future climate change selected for the SSP-RCP scenario framework. Accounting for this bias indicates RCP8.5 and other 'business-as-usual scenarios' consistent with high CO2 forcing from vast future coal combustion are exceptionally unlikely. Therefore, SSP5-RCP8.5 should not be a priority for future scientific research or a benchmark for policy studies." Justin Ritchie and Hadi Dowlatabadi, 'Why Do Climate Change Scenarios Return to Coal?', *Energy* 140 (1 December 2017):



Source: Matthew G. Burgess et al., 'IPCC Baseline Scenarios Have Over-Projected CO2 Emissions and Economic Growth', Environmental Research Letters 16, no. 1 (December 2020): Fig 5b., <https://doi.org/10.1088/1748-9326/abcdd2>

One concern might be that we discover a novel and cheap way to extract fossil fuels which could outcompete low carbon technology, in turn driving strong emissions growth. The historical trends suggest that this is unlikely. In the long-run, fossil fuel prices have been fairly stable, and have hovered around the same order of magnitude for more than a century. Way et al (2021) construct a probabilistic forecast using autoregression from past trends, which suggests that fossil fuel prices will be flat in the future.



Source: Way et al, 'Empirically grounded technology forecasts and the energy transition', Oxford Martin School, 2021, Fig. 3.

These forecasts account for large structural changes in fossil fuel extraction technology, such as the discovery of fracking.

The reason that fossil fuel prices will probably be fairly flat in the future is that there is a race between progress in extraction technology and use of the cheapest and most promising deposits. For example, in the early 2000s we discovered fracking technology, which allows us access to cheap gas. The best deposits are used up such that we then have to extract

from less promising and more expensive sites, so that the longer-term trend in natural gas prices is fairly flat.

The increase in coal use seems especially unlikely given the assumptions made in the SSPs about economic growth. For instance, in order for us to follow SSP5-RCP8.5, there would have to be very fast economic growth and technological progress, but meagre progress on low carbon technologies. This is implausible. In order to reproduce SSP5-8.5 with newer models, the models had to assume that average global income per person will rise to \$140,000 by 2100 and also that we would burn large amounts of coal.

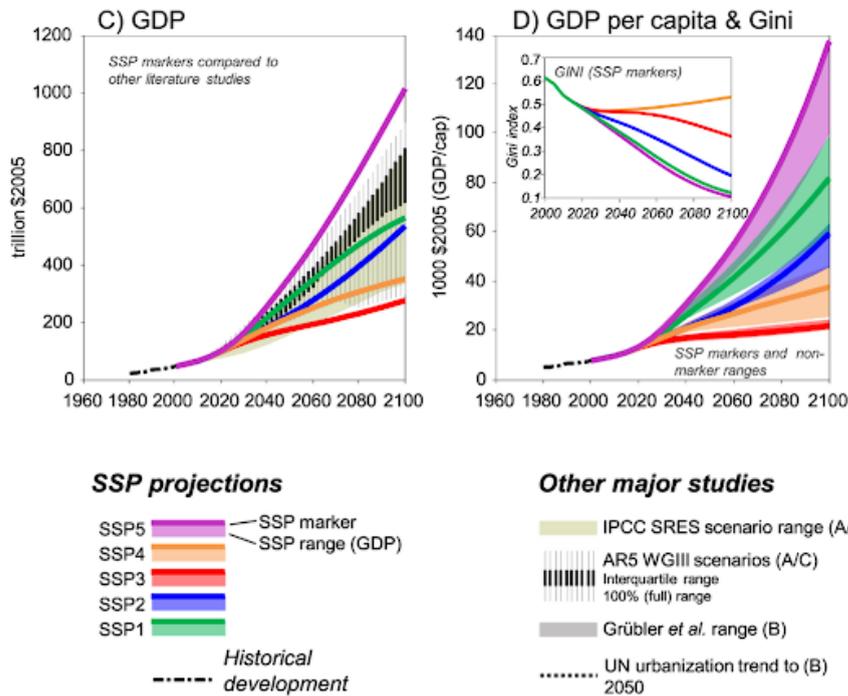


Fig. 2. Development of global population and education (A), urbanization (B), GDP (C), and GDP per capita and the Gini index (D). The inset in panel A gives the share of people without education at age of ≥ 15 years, and the inset in panel D denotes the development of the global (cross-national) Gini index. The SSPs are compared to ranges from other major studies in the literature, such as the IPCC AR5 (Clarke et al., 2014); IPCC SRES (Nakicenovic and Swart, 2000), UN, and Grubler et al. (2007). The colored areas for GDP (panel D) denote the range of alternative SSP GDP projections presented in this Special Issue (Dellink et al. (2016), Crespo Cuaresma (2016), Leimbach et al. (2016)).

Keywan Riahi et al., 'The Shared Socioeconomic Pathways and Their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview', Global Environmental Change 42 (1 January 2017): fig. 2.

It is difficult to imagine that in such a cornucopia, there would not also be a lot of progress on low carbon technology, and that countries would have greatly increased willingness to pay to protect the environment. The same argument applies to the other relatively high growth futures, such as SSP1 and SSP2.

1.2.3. Assumptions about economic growth

Economic growth is a key driver of emissions growth. The SSPs make assumptions about future economic growth that are crucial to future emissions projections. The historical track record of the projections of these models to date is poor and tends to overestimate both economic growth and emissions. The figure below shows the error in the Fifth Assessment Report (the 2013-14 IPCC report) and in the Shared Socioeconomic Pathways relative to observed growth:

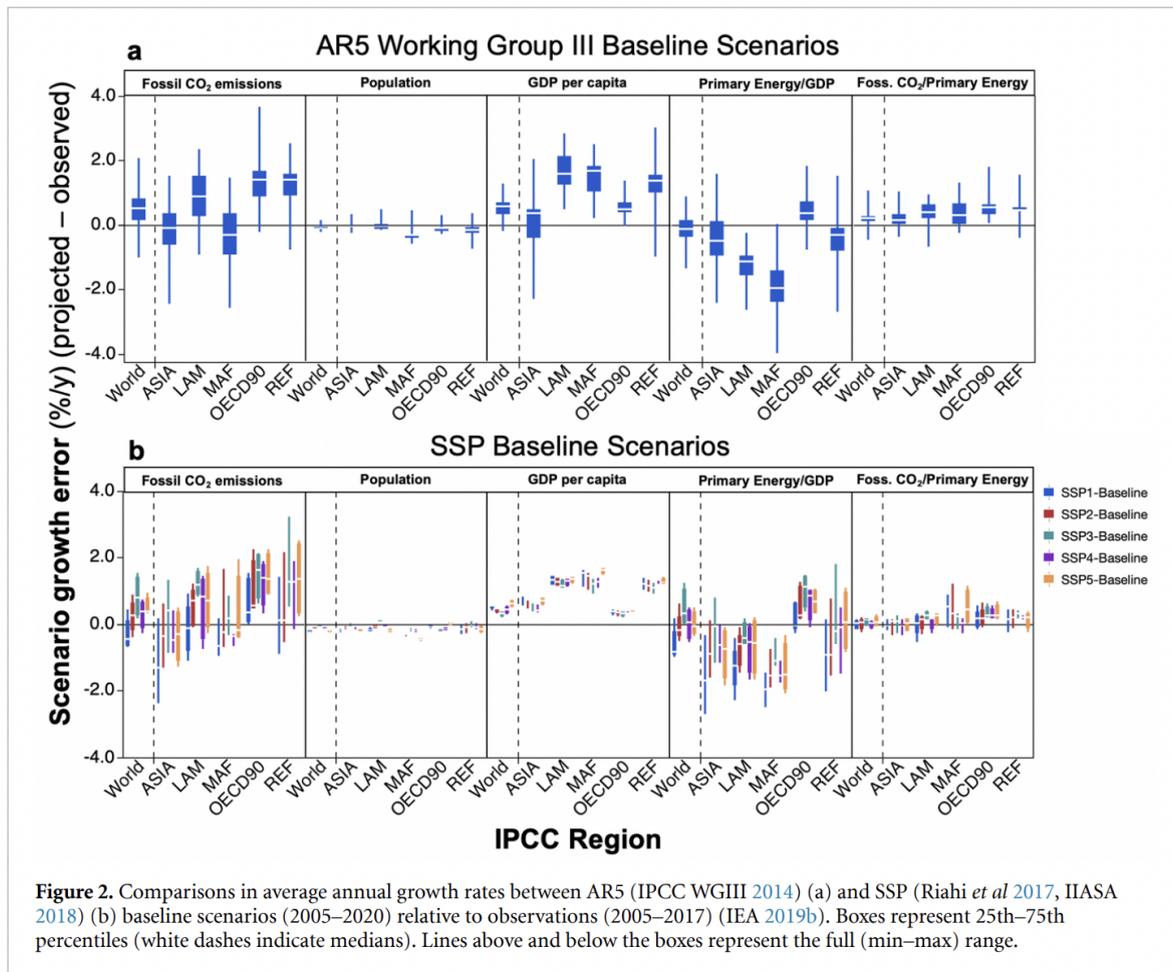


Figure 2. Comparisons in average annual growth rates between AR5 (IPCC WGIII 2014) (a) and SSP (Riahi *et al* 2017, IIASA 2018) (b) baseline scenarios (2005–2020) relative to observations (2005–2017) (IEA 2019b). Boxes represent 25th–75th percentiles (white dashes indicate medians). Lines above and below the boxes represent the full (min–max) range.

Source: Matthew G. Burgess *et al.*, ‘IPCC Baseline Scenarios Have Over-Projected CO₂ Emissions and Economic Growth’, *Environmental Research Letters* 16, no. 1 (December 2020): 014016, <https://doi.org/10.1088/1748-9326/abccd2>.

The error in projecting emissions growth is largely attributable to the overestimation of GDP per capita growth between 2005 and 2017.

These future projections of these models are also likely to be biased upwards. I discuss potential future growth trajectories in section 1.5.

1.3. Recent projections of emissions

Due to improving low carbon technology, strengthening climate policy, and the implausibility of a return to coal, recent models suggest that very high emissions scenarios are extremely unlikely and that we are most likely to follow a medium-low emissions pathway - around RCP4.5.⁹

⁹ For an overview, see Hausfather, ‘[Flattening the Curve of Future Emissions](#)’, Breakthrough Institute, 2021.

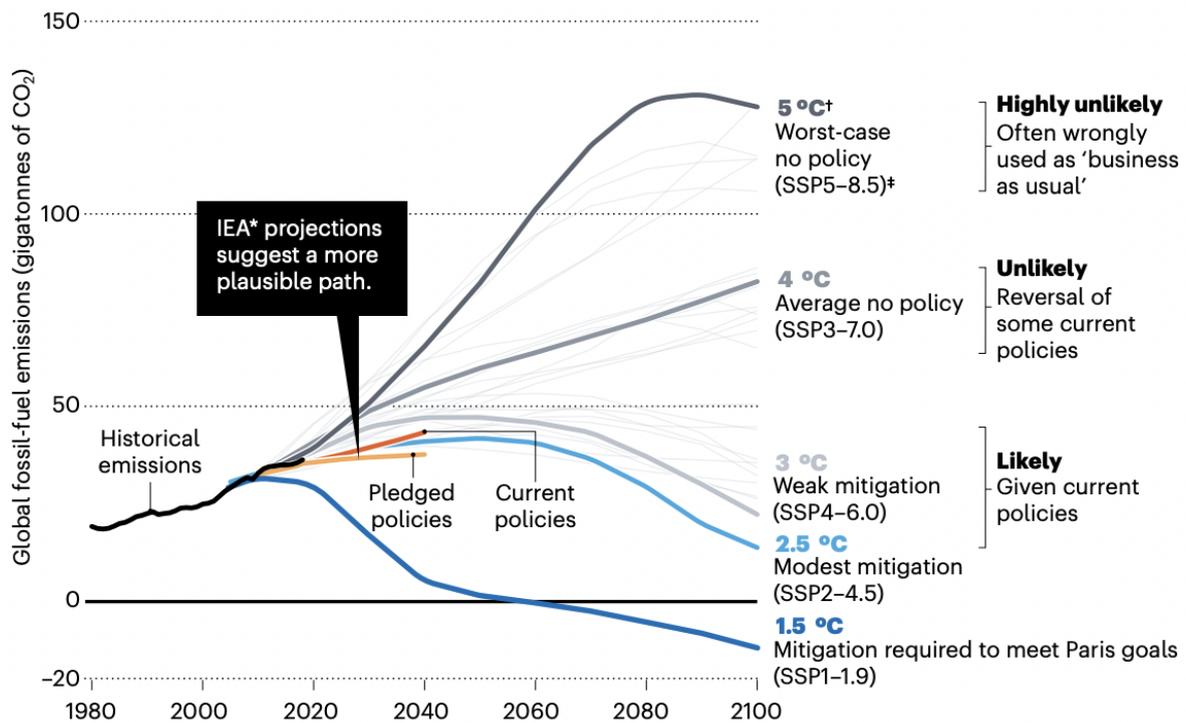
1.3.1. International Energy Agency

The International Energy Agency produces influential energy systems models that project likely future emissions, given certain socioeconomic, technological and climate policy assumptions. In 2015, the International Energy Agency projected that we were most likely to follow RCP6, on current policy.¹⁰ More recent estimates suggest that we are now most likely to follow RCP4.5, on current policy.

The chart below from Hausfather and Peters (2020) shows emissions on current policies and pledged policies, according to the International Energy Agency.

POSSIBLE FUTURES

The Intergovernmental Panel on Climate Change (IPCC) uses scenarios called pathways to explore possible changes in future energy use, greenhouse-gas emissions and temperature. These depend on which policies are enacted, where and when. In the upcoming IPCC Sixth Assessment Report, the new pathways (SSPs) must not be misused as previous pathways (RCPs) were. Business-as-usual emissions are unlikely to result in the worst-case scenario. More-plausible trajectories make better baselines for the huge policy push needed to keep global temperature rise below 1.5 °C.



*The International Energy Agency (IEA) maps out different energy-policy and investment choices. Estimated emissions are shown for its Current Policies Scenario and for its Stated Policies Scenario (includes countries' current policy pledges and targets). To be comparable with scenarios for the Shared Socioeconomic Pathways (SSPs), IEA scenarios were modified to include constant non-fossil-fuel emissions from industry in 2018.
 †Approximate global mean temperature rise by 2100 relative to pre-industrial levels.
 *SSP5-8.5 replaces Representative Concentration Pathway (RCP) 8.5.

Source: Hausfather and Peters, '[Emissions – the 'business as usual' story is misleading](#)', *Nature*, 2020.

The chart relies on IEA models of future energy systems. These are still probably too pessimistic on renewables. You can find the IEA's cost assumptions [here](#). They show the levelised cost of solar falling by 40% between now and 2030. But if historical trends continue, we should actually expect costs to decline by 89%. Trends may not continue, perhaps because we may be reaching saturation for renewables capacity additions, which

¹⁰ Wagner and Weitzman, *Climate Shock*, p. 31.

drive cost declines. But the default assumption should be that cost declines will continue. Going against this assumption has a poor track record, as noted above. The IEA revised down its short-term emissions projections in 2019, 2020 and 2021.¹¹

Moreover, climate policy will probably strengthen in the future. The IEA [retired](#) their current policies scenario after their 2019 report, arguing that the world was moving too quickly for a current policy scenario to be of much use. The latest IEA report projects emissions on near-term stated policies such as Paris commitments up to 2030 (STEPS), and the announced pledges scenario (APS 2021). Emissions on these scenarios are [lower](#) than RCP4.5.

Fossil CO2 emissions in IEA WEO and IPCC AR6 scenarios

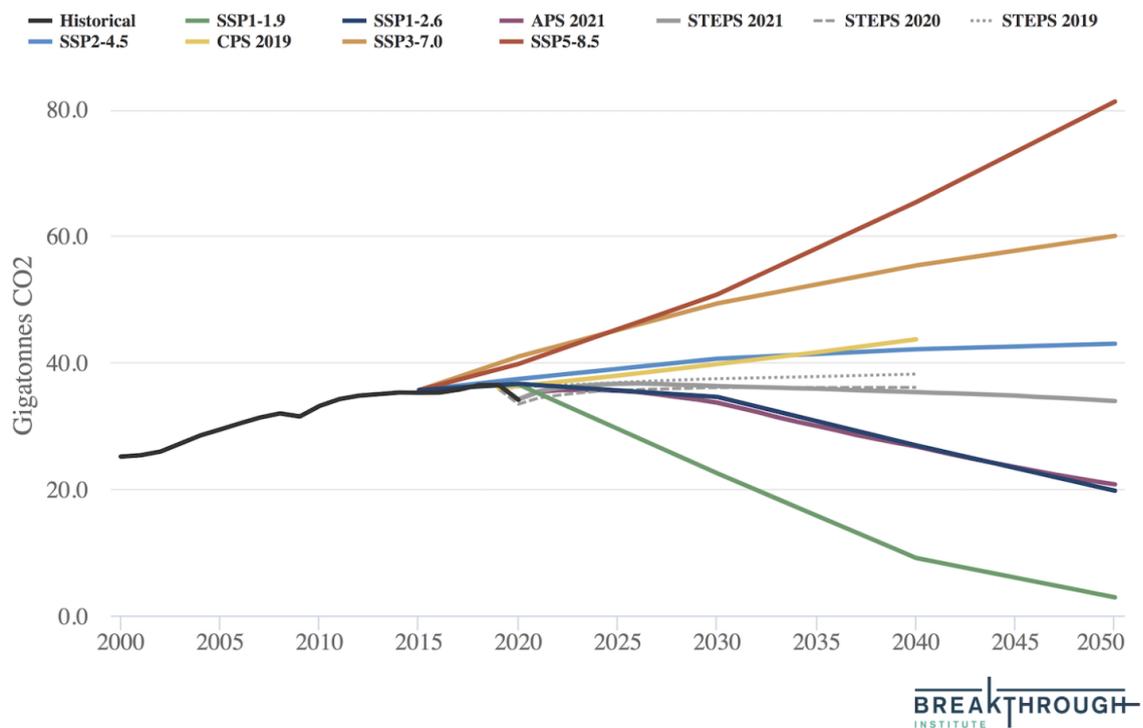
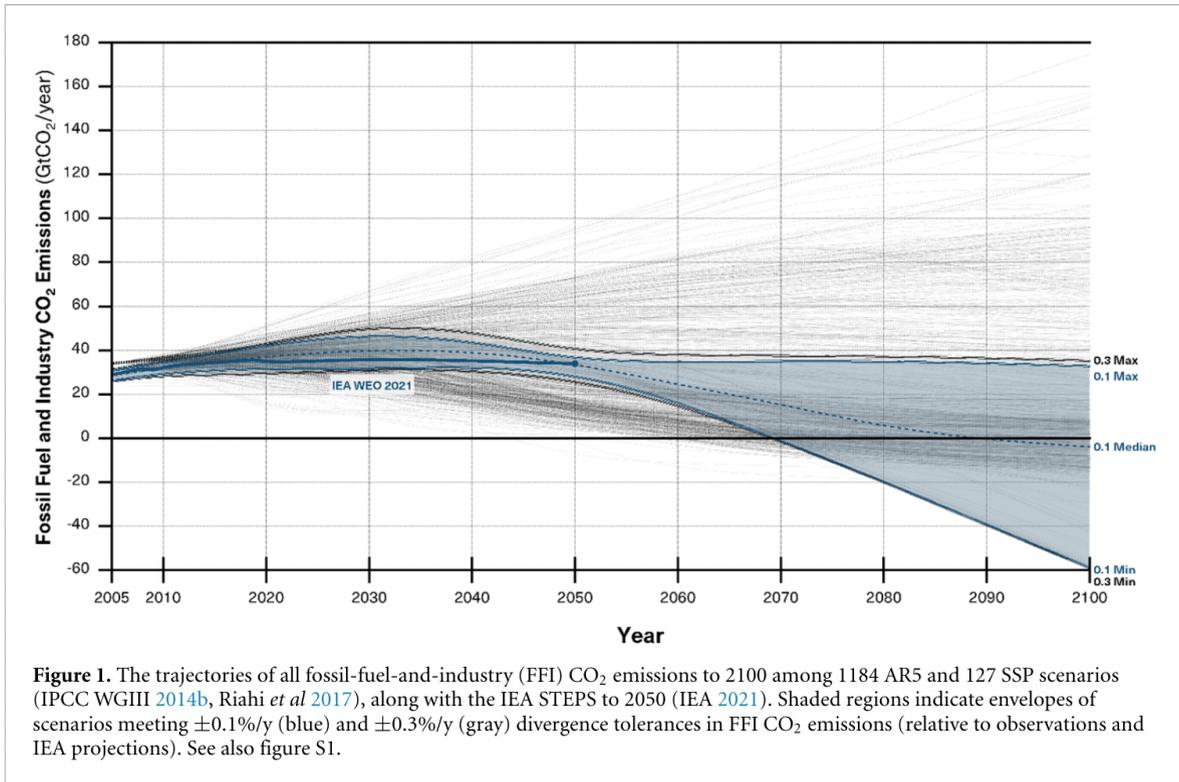


Figure 2: As with Figure 1, but including Shared Socioeconomic Pathways emissions scenarios from the [IIASA SSP Database](#) used in the [CMIP6 models](#) and featured in the [IPCC 6th Assessment Report](#).

Pielke Jr et al (2022) assess the plausibility of future emissions scenarios according to their compatibility with (1) historical fossil fuel emissions between 2005 and 2020 and (2) IEA projections to 2050. They rule out as implausible models that have a divergence of 0.3%/year on historical emissions and on IEA future emissions projections. Their results for future emissions are as follows:

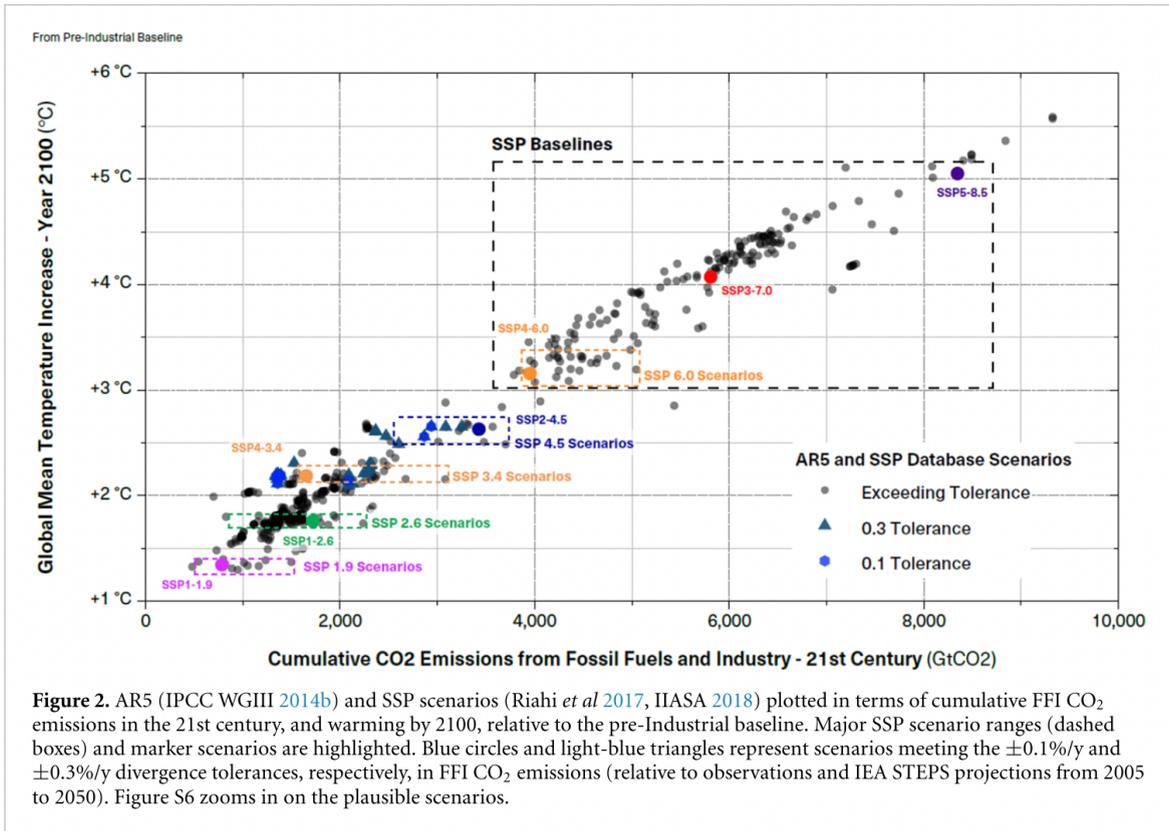
¹¹ “If IEA’s projections of near-term emissions continue to be revised down— as they have in recent years (IEA 2019, 2020, 2021)...” Roger Pielke Jr, Matthew G. Burgess, and Justin Ritchie, ‘Plausible 2005–2050 Emissions Scenarios Project between 2 °C and 3 °C of Warming by 2100’, *Environmental Research Letters* 17, no. 2 (February 2022): 3, <https://doi.org/10.1088/1748-9326/ac4ebf>.



Source: Roger Pielke Jr, Matthew G. Burgess, and Justin Ritchie, 'Plausible 2005–2050 Emissions Scenarios Project between 2 °C and 3 °C of Warming by 2100', *Environmental Research Letters* 17, no. 2 (February 2022): 3, <https://doi.org/10.1088/1748-9326/ac4ebf>.

As this shows, their 'best guess' scenario projects that we would decarbonise in around 2090.

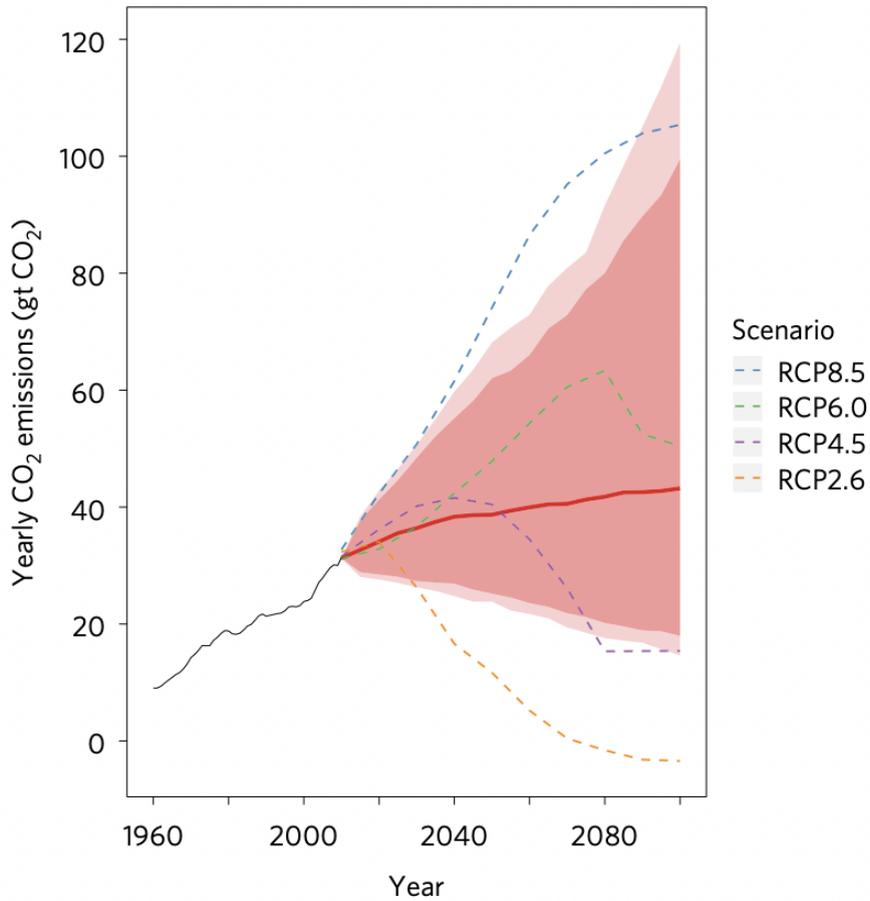
The chart below shows plausible scenarios according to Pielke Jnr. et al (2022) compared to SSP baselines. All of the SSP baselines are classed as implausible on their criteria; only those diagonally left of the SSP baseline rectangle are plausible.



Since IEA projections are likely too pessimistic on renewables, these classifications are probably overly pessimistic.

1.3.2. Liu and Raftery (2021)

In 2017, Raftery et al estimated the following probability distribution across emissions:



Source: Adrian E. Raftery et al., 'Less than 2 °C Warming by 2100 Unlikely', *Nature Climate Change* 7, no. 9 (September 2017): Fig 3a, <https://doi.org/10.1038/nclimate3352>.

This suggests that the most likely pathway is around RCP6 and that there is around 1-5% chance of RCP8.5. However, Liu and Raftery (2021) updated the estimates with more recent data up to 2015, which produced a more optimistic picture.

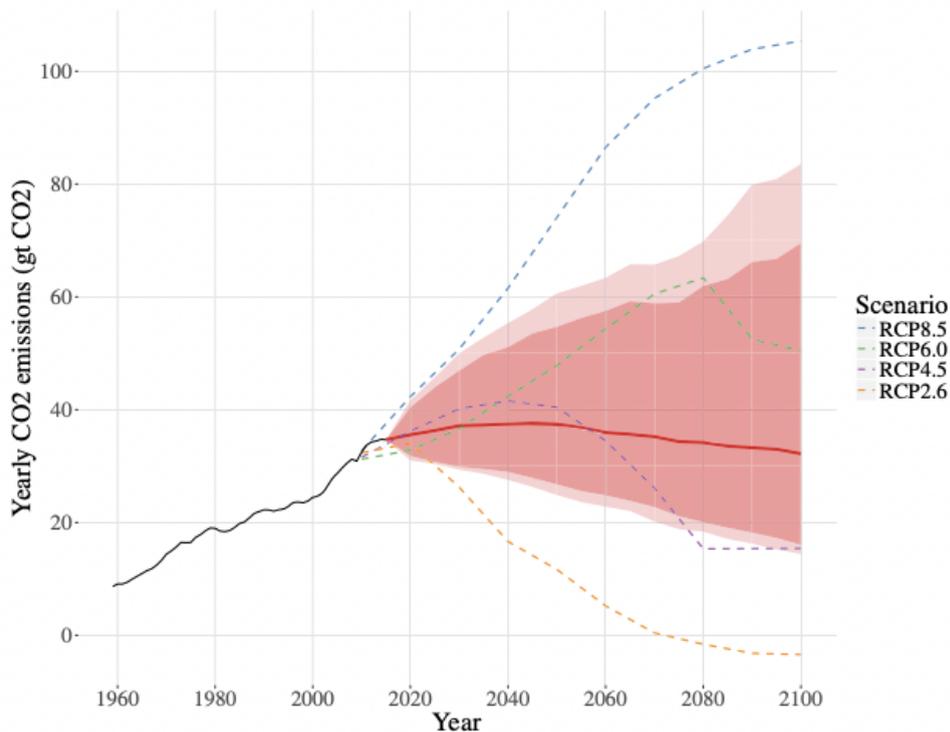
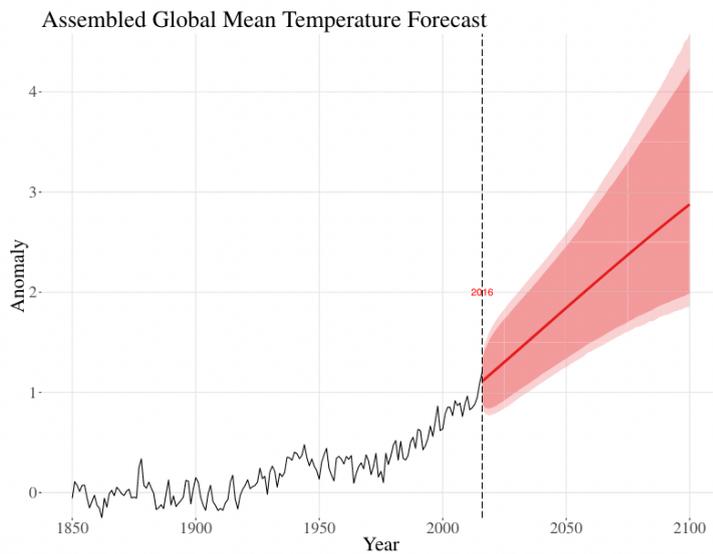


Fig. 1 Updated probabilistic forecast of CO₂ Emissions, based on data to 2015 and the method of Raftery et al.¹ The forecast median of yearly global emissions in 2100 is now 34 Giga tons.

Peiran R. Liu and Adrian E. Raftery, '[Country-Based Rate of Emissions Reductions Should Increase by 80% beyond Nationally Determined Contributions to Meet the 2 °C Target](#)', *Nature Communications Earth & Environment* 2, no. 1 (9 February 2021): Figure 1.

The most likely level of emissions is now between RCP4.5 and RCP6, and the chance of RCP8.5 is now well below 1%. I suspect that if their model were updated again using more recent data, projected emissions would be lower still.

Liu and Raftery et al (2021) combine their emissions estimates with the ensemble of CMIP5 climate models (which were used in the 2013-14 IPCC Fifth Assessment Report) to calculate the probability of different levels of warming:



Due to progress on emissions, likely warming estimated by Liu and Raftery (2021) has fallen relative to Raftery et al (2017). Liu and Raftery (2021) estimate that the 90% confidence interval spans from 2°C to 3.9°C. Relative to Raftery et al (2017), the median is 0.4°C lower and the 95th percentile is 1°C lower.¹²

These emissions estimates are produced in the following way. Raftery et al (2017) combines country-level estimates of future emissions using the Kaya identity.¹³

$$\text{Emissions} = \text{People} * \$ \text{ per person} * \text{emissions per } \$$$

They use UN population projections at the country-level. The model projects GDP per capita by assuming that other countries systematically converge towards the frontier level of GDP per capita, at a country-specific rate.

They model the logarithm of carbon intensity as following a linear trend, with randomness.¹⁴ At the global level, carbon intensity is actually declining linearly, but modelling the logarithm captures the idea that emissions reduction will get more difficult as we make more progress.

¹² “The median forecast for 2100 is 2.8 °C, with likely range (90% prediction interval) [2.1, 3.9] °C. The median is 0.4 °C lower than that of Raftery et al., the upper bound is 1.0 °C lower, while the lower bound is 0.1 °C higher. The tighter interval reflects the additional 5 years of data and the improved model.” Peiran R. Liu and Adrian E. Raftery, ‘Country-Based Rate of Emissions Reductions to Meet the 2 °C Target’, *Nature Communications Earth & Environment* 2, no. 1 (9 February 2021)

¹³ Adrian E. Raftery et al., ‘Less than 2 °C Warming by 2100 Unlikely’, *Nature Climate Change* 7, no. 9 (September 2017): 637–41, <https://doi.org/10.1038/nclimate3352>.

¹⁴ “We used the UN’s official 2015 population projections for all countries”; “There is a world frontier of GDP per capita, for which we use the United States as a proxy, and the GDP per capita of other countries converges to this world frontier at a country-specific rate.”; “We projected carbon intensity on the logarithmic scale for each country. We model the logarithm of carbon intensity as following a linear trend plus a first-order autoregressive process for each country.” Raftery et al., ‘Less than 2 °C Warming by 2100 Unlikely’.

The virtue of this estimate is that it is simple and more transparent than integrated assessment models, which often make use of opaque but implausible assumptions. It is also a great virtue that they give a 90% confidence interval, which is rarely done in the literature.

However, their approach has some drawbacks. Firstly, the assumption of independence between parameters is not true in practice and this likely systematically biases the estimate. Most importantly, their approach fails to capture the demographic transition - that higher economic growth will tend to lead to lower population growth. This biases their estimates upwards.

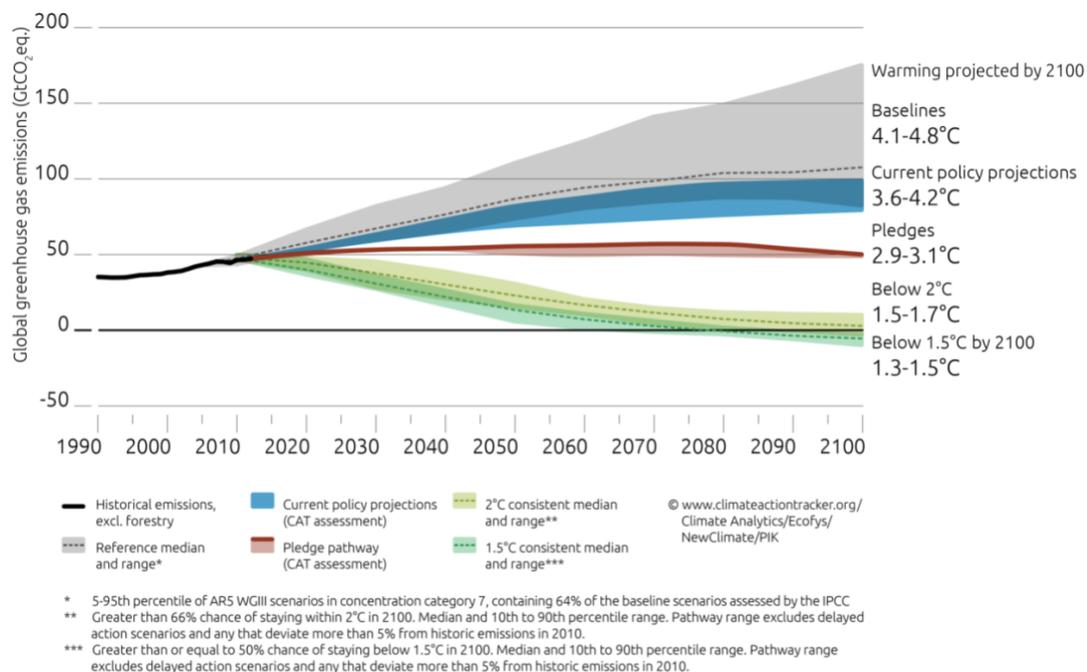
Secondly, Raftery et al use a 'assumption-light' approach to estimating the decline in emissions intensity that extrapolates the country-level trend in emissions intensity. The drawback of this approach is illustrated by the fact that they produced a different estimate with newer data: it turns out that the 2017 estimate was not a reliable estimate of future emissions. Sometimes, we have reasons to think that a trend will not continue into the future. Technological changes can bring non-linear changes in emissions, especially when the costs of key low carbon technologies are declining exponentially. Raftery et al's model does not capture such processes well.

Nonetheless, Liu and Raftery (2021) remains a very useful source on likely future emissions.

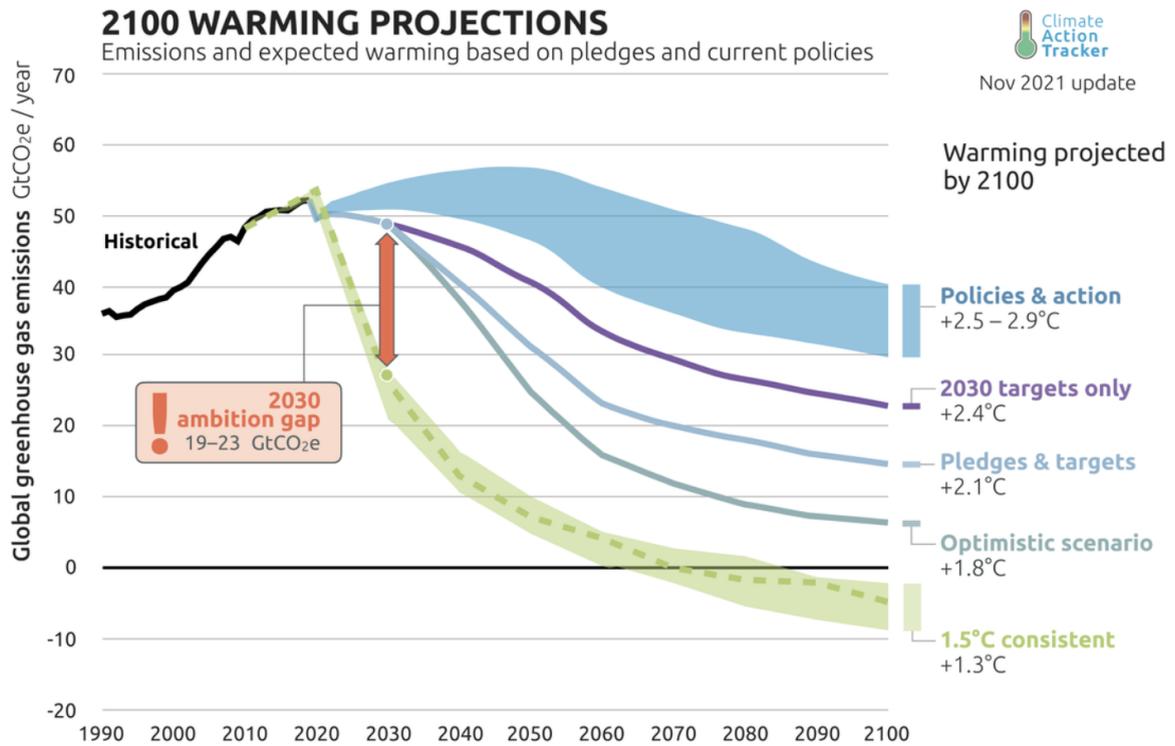
1.3.3. Climate Action Tracker

In 2015, Climate Action Tracker [estimated](#) that 4°C was the most likely scenario, on current policy.

Effect of current pledges and policies on global temperature



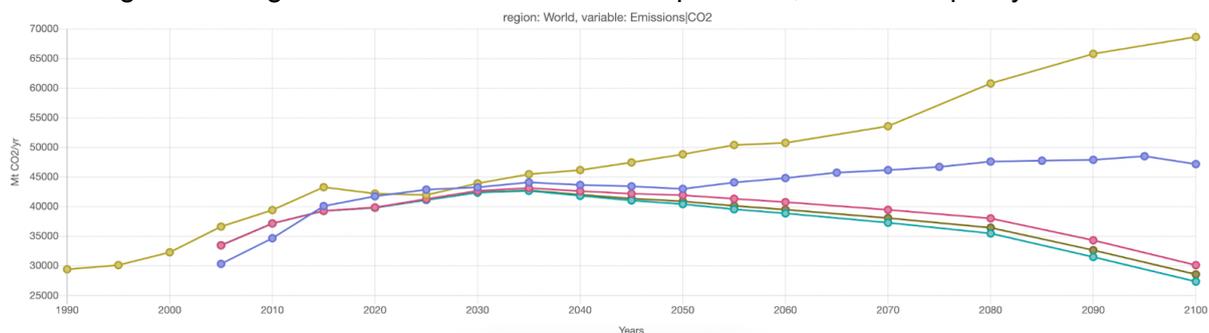
[Climate Action Tracker](#) now suggests that RCP4.5 is now the most likely scenario - this scenario implies around 2.7°C of warming.



As with the other studies, Climate Action Tracker finds marked progress on emissions.

1.3.4. Integrated Assessment Models

RCP8.5 was initially produced by Integrated Assessment Models called MESSAGE and REMIND. However, these and other Integrated Assessment models [now suggest](#) that something in the range of RCP4.5 to RCP7 is more plausible, on current policy.



It is difficult to assess how plausible such models are because they are complex and opaque. However, this does suggest that RCP8.5, on which emissions would reach 120 Gt CO₂ per year by 2100 now looks very unlikely. Moreover, as discussed above, these models are probably too pessimistic about cost declines in renewables and batteries, and many of them seem to rest on erroneous assumptions about economic growth and coal consumption.

1.3.5. Meinshausen et al (2020)

For a long time, pledges made following the Paris Agreement were insufficient to meet the Paris goals of limiting warming to 2°C.¹⁵ However, due to recent progress on climate policy, if current pledges are met, we have a better than 50% chance of limiting warming to 2°C and that the chance of RCP4.5 is less than 5%.

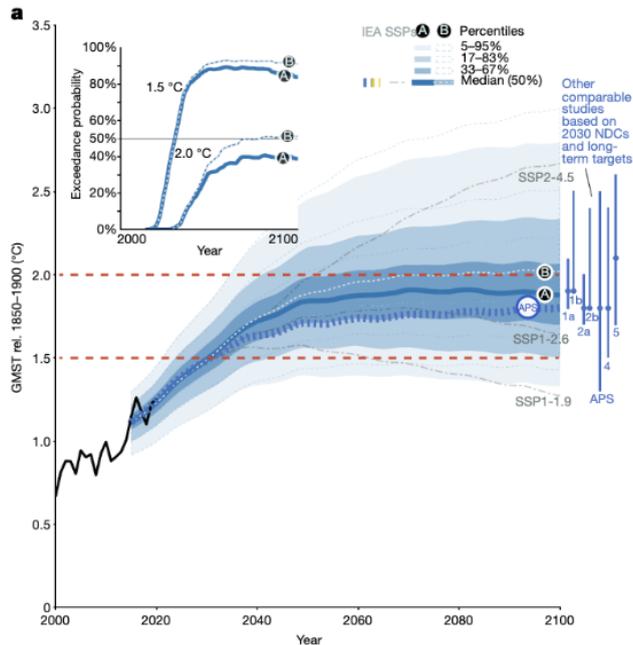


Fig. 3 | Global mean temperature projections based on 2030 NDCs show a wide range, whereas those based also on long-term targets stay just around or below 2 °C—with limited additional effects by the GMP. a, Warming due to NDCs and long-term targets. Global mean temperature projections based on all officially submitted NDCs and long-term targets as of the end of COP26 in addition to the announcement by India (1 November 2021), with a sensitivity case considering full implementation of NDC targets (A) and only unconditional targets (B). Inset, scenario A has a better than 50% chance and scenario B has a roughly 50% chance, but neither scenario has a likely 67% chance to stay below 2 °C. **b,** Greenhouse gas emissions without LULUCF for the considered pathways extending 2030 NDCs (orange) and the two sensitivity

Source: Malte Meinshausen et al., ‘Realization of Paris Agreement Pledges May Limit Warming Just below 2° C’, *Nature* 604, no. 7905 (2022): fig. 3a.

There seems a decent chance that many countries will not in fact meet their pledges. But this does illustrate how climate ambition has increased over the last decade.

1.3.6. Summary

A range of recent studies suggest that something around RCP4.5 is now the most likely emissions scenario, on current policy. Given that policy is likely to strengthen in the future, the most likely scenario seems likely to be between RCP2.6 and RCP4.5.

It is much less clear what the risk of much higher emissions is, but the high emissions SSP5-8.5 seems extremely unlikely (I think closer to 1 in 1,000 than 1 in 100) given the

¹⁵ Joeri Rogelj et al., ‘Paris Agreement Climate Proposals Need a Boost to Keep Warming Well below 2 °C’, *Nature* 534, no. 7609 (30 June 2016): 631–39, <https://doi.org/10.1038/nature18307>.

required increase in coal use per person. One qualification to this might be that progress in AI could drive an explosion in economic growth and in carbon emissions. I discuss this below.

1.4. Will renewables and batteries solve climate change on their own?

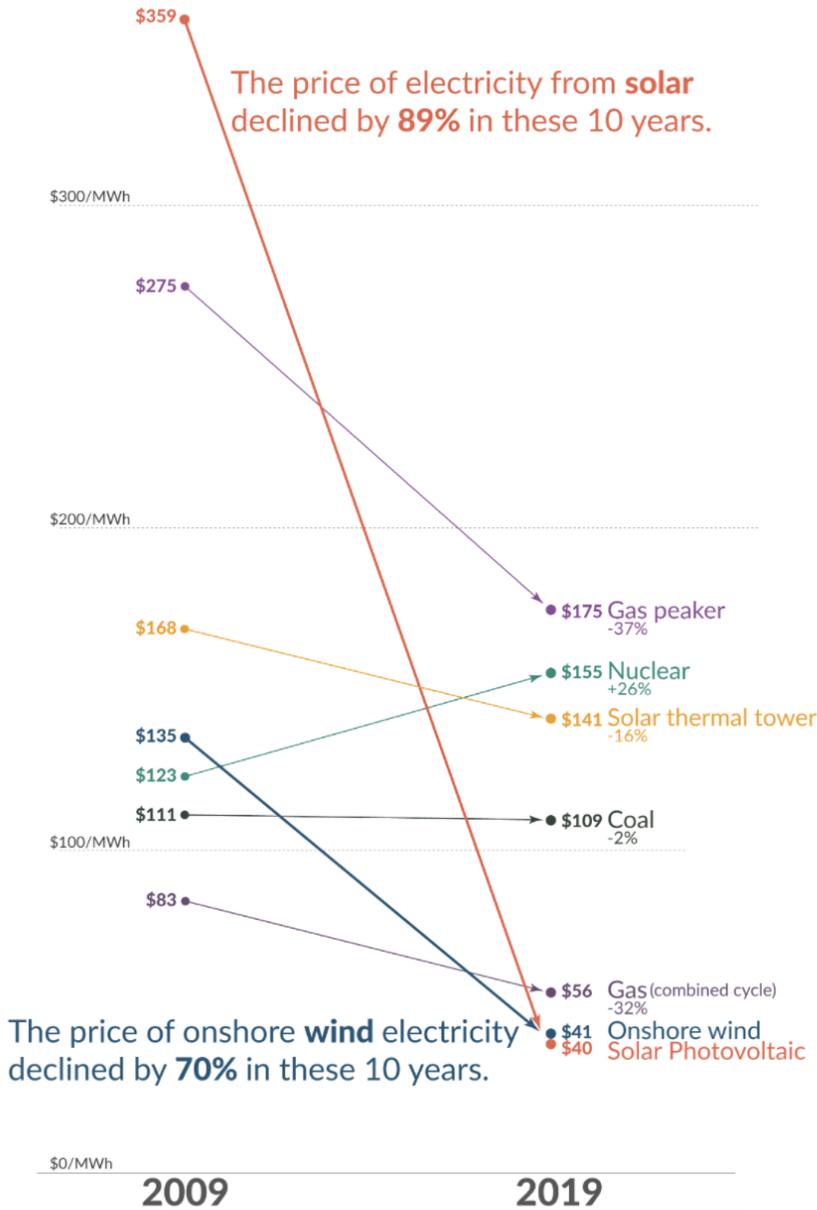
I have argued that we now have more cause for optimism on emissions than we have had for some time. However, some argue that even this view is too pessimistic because it underestimates likely progress in renewables.

There are two versions of this argument. The more naive version argues that solar, wind and batteries are now cheaper than fossil fuels and so will soon replace all fossil fuel infrastructure, even without a strong change in climate policy. It might be natural to draw this conclusion from this chart, for example.

The price of electricity from new power plants



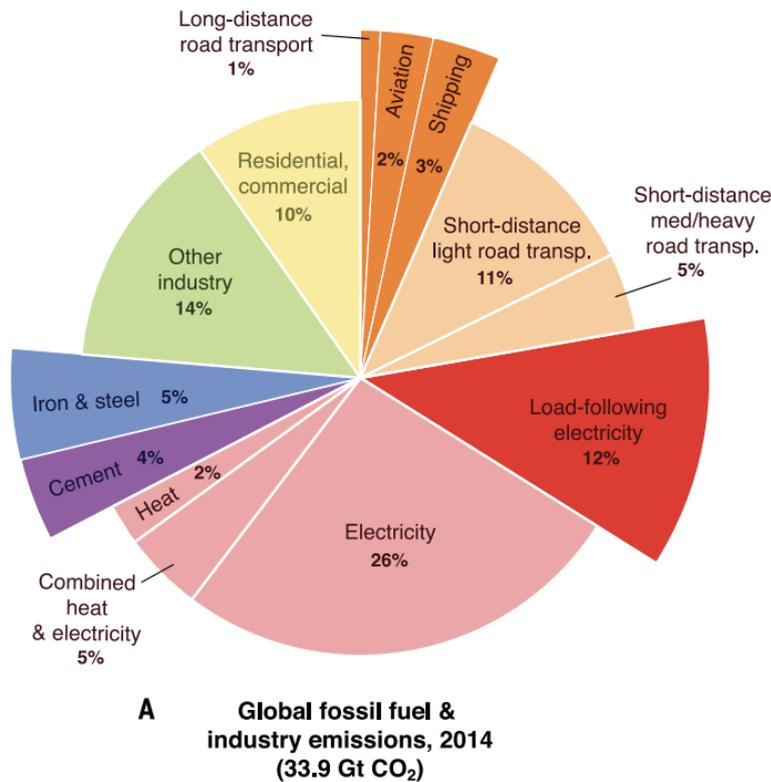
Electricity prices are expressed in 'levelized costs of energy' (LCOE). LCOE captures the cost of building the power plant itself as well as the ongoing costs for fuel and operating the power plant over its lifetime.



Data: Lazard Levelized Cost of Energy Analysis, Version 13.0
 OurWorldinData.org - Research and data to make progress against the world's largest problems. Licensed under CC-BY by the author Max Roser.

There are several problems with this line of argument. Firstly, the levelised cost metric used in the chart above is a misleading way to compare different energy technologies, especially intermittent and non-intermittent sources. Because solar and wind are intermittent, increasing the penetration of renewables imposes costs on the system in terms of additional transmission, storage and backup generation. To understand whether solar and wind are going to take over the whole electricity system, we need to understand the costs of the whole system once all of the apparatus has transitioned, not the costs to a businessperson of building a marginal solar plant.

Secondly, electricity is only around 40% of emissions from fossil fuels and industry. A substantial fraction of emissions come from other sectors like industry and transport. Some of these sectors can be electrified at reasonable cost, as with electric cars, but others, such as aviation, shipping, steel and cement are much more difficult to electrify. The chart below shows 'difficult-to-eliminate' emissions.



Source: Steven J. Davis et al., 'Net-Zero Emissions Energy Systems', *Science* 360, no. 6396 (29 June 2018): fig. 2, <https://doi.org/10.1126/science.aas9793>.

So, even if renewables take over a substantial chunk of electricity, we would still have a long way to go to get to net zero emissions.

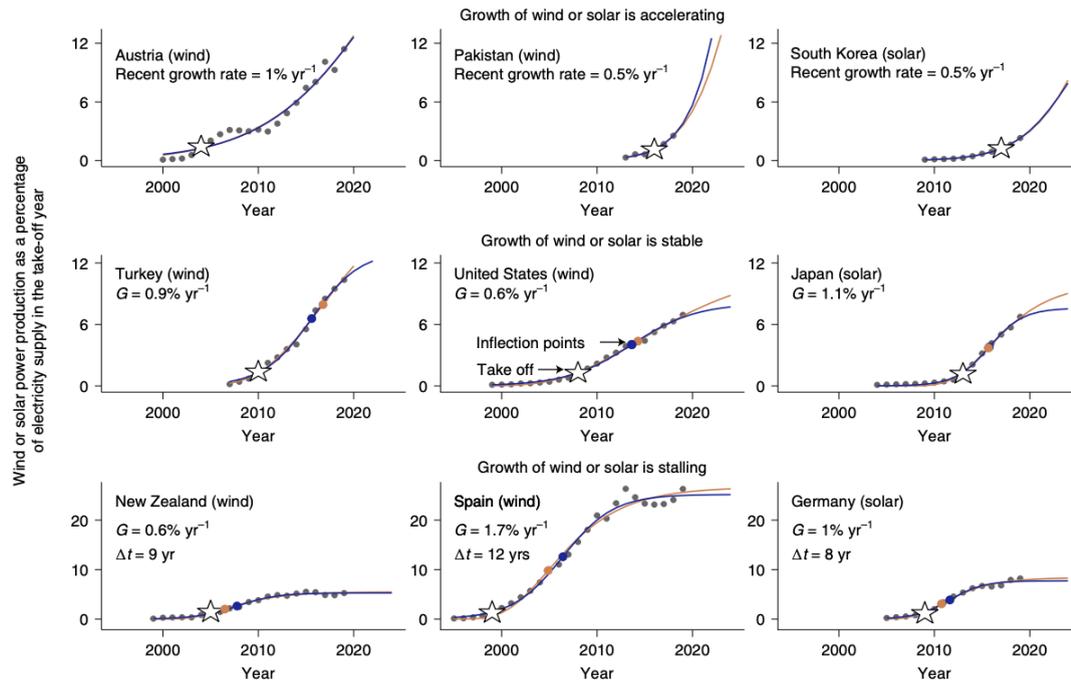
The most sophisticated version of the 'renewables takeover' argument has recently been presented in a wonderful paper by [Way et al \(2021\)](#). They argue that if exponential progress in renewables, batteries and hydrogen (or other power-to-X technology) continues, we will nearly completely decarbonise the economy by 2040 *and* save trillions of dollars, even without accounting for the costs of climate change.

The main drawback of this paper is that it is not clear whether the cost declines in these low carbon technologies will continue, and whether, even if they do, they will be scaled up sufficiently. So far, the cost declines have been driven by *Wright's Law*, which predicts that costs drop as a power law of cumulative production.¹⁶ This relationship is also called an *experience curve* or *learning curve*. According to Wright's Law, doubling cumulative capacity

¹⁶ Way et al, 'Empirically grounded technology forecasts and the energy transition', Oxford Martin School, 2021, 4-5.

leads to a fixed percentage decline in cost. This means that to produce a given percentage decline in cost, we need to add more and more capacity in absolute terms.

Unfortunately, there are social and political barriers to the dramatic scale-up of renewables. In many countries, even as costs decline exponentially, renewables deployment is starting to slow:



Source: Aleh Cherp et al., 'National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets', *Nature Energy* 6, no. 7 (July 2021): Fig. 3, <https://doi.org/10.1038/s41560-021-00863-0>.

In Cherp et al's sample of 27 countries, the growth of renewables is only accelerating in 5 countries, is stable in 11 and is stalling in 11.¹⁷ As Cherp et al (2021) note,

“Declining costs of technologies have already led to a relatively high growth in the OECD, although currently this growth is becoming constrained on sociotechnical and political rather than economic grounds”.¹⁸

One important barrier is land use concerns. Solar and wind require much greater land area than fossil fuels or nuclear power to produce a given amount of energy. So, there is more scope for local opposition. The chart below shows the land area that would be required to replace South Korea's oil consumption with hydrogen generated from solar and wind, compared to nuclear power (advanced heat sources).

¹⁷ Aleh Cherp et al., 'National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets', *Nature Energy* 6, no. 7 (July 2021): Table 2, <https://doi.org/10.1038/s41560-021-00863-0>.

¹⁸ Aleh Cherp et al., 'National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets', *Nature Energy* 6, no. 7 (July 2021): p. 751, <https://doi.org/10.1038/s41560-021-00863-0>.



Each colored outline represents the total area that would be required for the siting of each type of resource if it were to be the only one used to generate enough hydrogen to replace current oil consumption in South Korea.

Source: Lucid Catalyst and TerraPraxis, 'Missing Link to a Livable Climate', 2020, p. 34.

Note that this only accounts for oil and not for coal and gas consumption. Restrictive land use policies are a well-known problem in rich countries.¹⁹ This seems like a significant barrier to the scale-up of renewables.

Overall, while progress in renewables provides cause for optimism, it would be too hasty to conclude that renewables will soon completely solve climate change.

1.5. Worst-case emissions scenarios

A range of models now suggest that something around RCP4.5 is the most likely scenario, on current policy. However, much less attention has been devoted to understanding the probability of worst-case scenarios, which may be disproportionately important for climate risk.

¹⁹ Edward L. Glaeser, Joseph Gyourko, and Raven Saks, 'Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices', *The Journal of Law and Economics* 48, no. 2 (1 October 2005): 331–69, <https://doi.org/10.1086/429979>.

1.5.1. How high could fossil fuel emissions be?

If climate policy went badly wrong, how much fossil fuel could we burn, and how much carbon could we release into the atmosphere? This question has received surprisingly little attention in the literature.

Fossil fuel reserves and resources are defined as follows:

Reserves = A fossil fuel deposit that is economically exploitable at today's prices and using today's technology.

Resources = A proven fossil fuel deposit that cannot currently be exploited or an unproven but geologically possible fossil fuel deposit that may be exploitable in the future.²⁰

Reserves are dynamic: they change with technology and with market prices. If the price of natural gas rises, then it may become economically viable to exploit previously non-viable deposits.

Estimates of resources also change over time as our knowledge of particular fossil fuel deposits improves. There is much more uncertainty about total resources than there is about reserves as there are much stronger market incentives to understand remaining reserves and industrial actors naturally have much better knowledge of deposits that they are currently exploiting than of potentially exploitable deposits.

Moreover, the definition of resources is not completely clear. What does it mean to say that a deposit 'may be exploitable' in the future? The most natural interpretation is that the deposit would become economically exploitable given possible changes in technology and prices of fossil fuels. But the claim that a change is 'possible' is extremely broad; such a change would not be very action-relevant if it is extremely unlikely. It is difficult to know which technological and price changes would be needed if we were to try to exploit a substantial fraction of the remaining fossil fuel resources. For example, one researcher told me that coal seams under the North Sea are classed as part of Britain's coal resource, but it is extremely unlikely that these resources will ever be extracted, though we could if we really wanted to.

The IPCC uses the following estimates of fossil fuel reserves and resources from the German organisation BGR.²¹

	Billion tonnes of carbon
Fossil fuel reserves	900
Fossil fuel resources	12,360

²⁰ BGR, 'Energy Study: Data and Developments Concerning German and Global Energy Supplies' (Hannover, 2019), 192.

²¹ BGR, 'Energy Study: Data and Developments Concerning German and Global Energy Supplies' (Hannover, 2019), Fig. 5.12.

I have created a sheet which summarises some data on fossil fuel reserves and resources [here](#).

For reference, we have [released](#) 460 billion tonnes of carbon since the Industrial Revolution. Coal accounts for 63% of remaining fossil reserves and 93% of remaining resources ([fossil fuel sheet](#)).

However, a crucial point that is often not made clear in the literature is that *not all resources are recoverable*. An underground coal seam may be classed as a resource even if we might only expect to extract a fraction of the coal from this deposit. Ritchie and Dowlatabadi claim that recovery rates for surface coal resources are around 80%, but for underground coal seams, they would typically be 50% and for some seams as low as 20%.²² The IPCC's Fifth Assessment Report, which produced a similar estimate of remaining fossil fuel resources to the Sixth Assessment, relied on Rogner et al (2012), which says "Resources are shown as in situ amounts; the eventually extractable quantities will be significantly lower". Rogner et al (2012) contend that a lot of coal resources are very hard to extract: they are in narrow seams more than a kilometre below the surface. Extraction rates in geologically difficult deposits are typically below 40%, and with current technology would only be around 20%.²³

This point is often misunderstood in the literature, which is understandable given the lack of clarity about the meaning of 'resource' in the IPCC reports. For example, Winkelmann et al (2015) model the effects of releasing up to 12 trillion tonnes of carbon into the atmosphere and says "We hereby cover the full range of available carbon resources".²⁴

When we are evaluating a worst-case scenario, we actually need to estimate not total resources, but rather what is known as the 'Ultimately Recoverable Resource'.

(Note that, confusingly, only a fraction of coal *reserves* are recoverable, while close to 100% of oil and gas reserves are recoverable.)²⁵

²² "Common recovery factors are 80 percent for strip mining and 50 percent for underground reserves. However, Zimmerman (1983) studies the US coal industry and suggests that 50 percent is a reasonable estimate for reserves recoverable from any deposit. Reserve recovery in some regions such as India can be as low as 20% (Bauer et al. 2016)." Justin Ritchie and Hadi Dowlatabadi, 'The 1000 GtC Coal Question: Are Cases of Vastly Expanded Future Coal Combustion Still Plausible?', *Energy Economics* 65 (1 June 2017): 9 n2, <https://doi.org/10.1016/j.eneco.2017.04.015>.

²³ "While coal resources are 20 times higher than known extractable coal reserves, there is uncertainty about the minable portion of these in situ quantities. Information on the geomining conditions of coal resources is insufficient for a reliable production assessment. For example, most of the better delineated coal resources are situated at greater depths and thus belong to geomining categories CIII and CIV. Extraction ratios in geologically difficult coal deposits can be below 40% (Kundel, 1985 ; Daul, 1995 ; USGS, 2009). Since many of the 'in situ' hard-coal resource deposits are in narrow seams at depths of more than 1000 m, an overall recovery rate of 20% may well be achievable practically. For example, 60% of coal resources in China are found at depths deeper than of 1000 m (Pan, 2005 ; Minchener, 2007). Without new extraction methods, a 20% recovery rate puts the portion of coal resources that eventually could become available as reserves to 87,154 EJ." H-H. Rogner et al., 'Energy Resources and Potentials', in *Global Energy Assessment - Toward a Sustainable Future* (Cambridge University Press, 2012), 464.

²⁴ Ricarda Winkelmann et al., 'Combustion of Available Fossil Fuel Resources Sufficient to Eliminate the Antarctic Ice Sheet', *Science Advances* 1, no. 8 (1 September 2015): e1500589, <https://doi.org/10.1126/sciadv.1500589>.

²⁵ Conversation with Sandro Schmidt, 21 October 2021.

The same conceptual problems about the definition of resources infect the definition of Ultimately Recoverable Resources. Because the definition of a resource is unclear, it is hard to understand the technology assumptions that determine ultimately recoverable resource estimates. Rather than use these sorts of imprecise definitions, it would be preferable either to (1) provide a probability distribution across amounts of possible future coal extraction, or (2) outline the exact conditions in which a resource is likely to be extracted.

Will we try and extract all of the coal?

Ritchie and Dowlatabadi (2017) argue that much of the remaining coal resources are unlikely to ever be extracted. A large fraction of the remaining coal resources are underground deposits that could only feasibly be extracted using a technology called 'underground coal gasification'. This involves drilling injection wells into a coal seam, and igniting the seam so that temperatures reach 500 to 900°C. This converts the coal to a mixture of CO₂, CO, CH₄ and H₂, which is removed using extraction wells drilled into the seam.²⁶ The CH₄ can then be burned to produce energy.

Ritchie and Dowlatabadi (2017) are pessimistic about the prospects for underground coal gasification. Underground coal gasification was first proposed in the 19th century and has failed the test of commercialisation for more than a century, with trials only having been run for a few days or weeks. They argue that challenges around siting and environmental costs make the future prospects look dim:

“Experience from UCG test projects have indicated significant constraints on site selection. For example, a 1997 pilot in Spain at a depth of 600 meters highlighted the importance of avoiding aquifer systems because of the potential for explosions. In this case, geological subsidence shifted the underground structure, leading to collapse and a subsequent explosion (Walker 2007). UCG pilots in many locations have caused severe groundwater contamination that persists for years after gasification ceased, with high concentrations of phenols and PAHs readily detected in aquifers extending dozens of kilometers from the gasification site (Campbell et al. 1979; Friedmann et al. 2009; Klimenko 2009; Liu et al. 2007). Given the documented public response to large-scale coal synfuel and syngas projects (Yanarella and Green 1987), it reasonable to expect that any social license for operation of UCG facilities will face significant opposition, even if many of its environmental challenges are successfully addressed...

Despite more than a century of experimentation, recent meta-assessments conclude that UCG still needs decades of foundational research to establish any reasonable estimate of its commercial potential (Couch 2009). In this context UCG would need to contend with rapid progress in commercial-scale renewable energy, unconventional oil and gas and more energy efficient technologies (IEA 2016). This is a challenging environment to justify a new wave of sustained public or private funding for research and development of UCG. Thus, any plausible future reference case for global coal recovery should not include estimates of total resources which are implicitly or

²⁶ Ritchie and Dowlatabadi, 'The 1000 GtC Coal Question', p. 18.

explicitly consistent with theoretical potentials of UCG or other similar hypothetical technologies.”²⁷

Other coal resource experts I have spoken to also thought that underground coal gasification would play at most a marginal role in the future.

The vast majority of remaining coal resources are [located](#) in five countries: the US, China, Russia, Australia, and India. So, the amount of coal extracted will depend in part on the likely environmental restrictions in these countries. Environmental regulations in China and the US are increasing. Coal extraction remains an extremely important political issue in Australia, and it is widely thought that the Labour Party lost the last election for being insufficiently pro-coal. Australian coal is high quality and in high demand in neighbouring growing Asian countries.

Assumptions about technological development in fossil fuel extraction technology need to be consistent with assumptions about technological development in other domains. Worlds in which we develop advanced coal extraction technology are also likely to be worlds in which technological progress in other domains, including low carbon technology, will be strong. As Ritchie and Dowlatabadi (2017) note:

“if it is appropriate to consider the implications and recovery rates of coal consistent with UCG deployment in reference global energy scenarios, it would be equally appropriate to consider the role of experimental technologies such as nuclear fusion.”²⁸

One natural concern with these arguments is that we may not have a good sense of the fossil fuel extraction technologies that we will develop in the future.

Ultimately recoverable resources

There are surprisingly few estimates of ultimately recoverable resources in the literature, and the data on coal seems to be of much worse quality than the data for oil and gas. As Ritchie and Dowlatabadi (2017) note

“However, efforts to determine the potentially recoverable portion of world coal resources have been fragmented, compromising time-series analyses with notoriously inconsistent and poor data.”²⁹

Estimates of ultimately recoverable fossil fuel resources are show in the table below:

Study	Billion tonnes of carbon
Ritchie and Dowlatabadi (2017) for coal + author calcs using BGR data	1,200
Welsby et al	2,860

²⁷ Ritchie and Dowlatabadi, ‘The 1000 GtC Coal Question’, pp. 19-20.

²⁸ Ritchie and Dowlatabadi, ‘The 1000 GtC Coal Question’, p. 20.

²⁹ Ritchie and Dowlatabadi, ‘The 1000 GtC Coal Question’, p. 3.

Mohr et al low estimate	1,040
Mohr et al best guess	1,580
Mohr et al high estimate	2,500

These are summarised in the [fossil fuel data sheet](#).

I don't have a great deal of confidence in these numbers. I spoke to Steve Mohr about how they produced their estimates on coal resources and he suggested that a lot of judgement calls are involved. As I understand it, a lot of the knowledge required is highly mine-specific, requiring in-depth knowledge from in-country geologists, which is not publicly shared or transparent.

In any case, all of these estimates are lower than the emissions implied in some emissions scenarios considered in the literature. On SSP5-8.5, we would burn an additional 2,200 billion tonnes of carbon by 2100, which exceeds some of the estimates above.³⁰ Climate scientists have also outlined an extended version of SSP5-8.5, on which emissions continue to 2500 and we burn 5 trillion tonnes of carbon, which exceeds all of the estimates shown above.³¹

I will assume that ultimately recoverable resources in a worst-case scenario total 3,000 billion tonnes of carbon, which is higher than all of the estimates.

For comparison:

- Fossil fuel and cement emissions since 1750 = 464 billion tonnes of carbon.
- Fossil fuel and cement emissions in 2018 = 10 billion tonnes.
- On RCP4.5, emissions 2019 to 2100 = 848 billion tonnes.³²

If emissions remained at current levels, it would take 300 years to burn through all of the ultimately recoverable fossil fuel resources.

How plausible is it that we will burn through all of the fossil fuel resources? I think there are two possible scenarios on which this could happen.

³⁰ IPCC, Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change (UNFCCC, 2021), SPM.7.

³¹ Katarzyna B. Tokarska et al., 'The Climate Response to Five Trillion Tonnes of Carbon', *Nature Climate Change* 6, no. 9 (September 2016): 851–55, <https://doi.org/10.1038/nclimate3036>; Malte Meinshausen et al., 'The RCP Greenhouse Gas Concentrations and Their Extensions from 1765 to 2300', *Climatic Change* 109, no. 1–2 (1 November 2011): 213, <https://doi.org/10.1007/s10584-011-0156-z>.

³² IPCC, Sixth Assessment Report: Working Group I The Physical Science Basis, SPM.7.

1.5.2. Explosive economic growth

Economic growth is an important driver of emissions growth. The SSPs make a range of assumptions about economic growth. On the high growth future, SSP5-8.5, global GDP increases by upwards of 2% per year.

Is it possible that growth could be much higher than this? If we had asked an economist in 1800 how fast growth would have been in the next 200 years, none would have predicted that it would have risen to 1% or 2%. This was completely unprecedented, but nevertheless occurred. Prior to that, for thousands of years, living standards barely improved and the growth rate was probably only 0.1%. The rate of economic growth rate is itself growing over time. If we think that this long-term trend will continue in the future, then we should expect the growth rate itself to increase again in the next few hundred years.

A second related reason to think that there might be explosive growth at some point in the next 100 years is progress in Artificial Intelligence. According to endogenous growth theories, growth is caused by *innovation*: the production of new ideas and new technologies. On this theory, population growth is the primary driver of economic growth: the more people there are, the more ideas there are and the more innovation there will be.

Many people expect that at some point this century, AI systems will become so advanced that we will be able to automate innovation. Quite when AI systems might reach the human level at innovation remains unclear. The most sophisticated attempt to forecast transformative AI is by Ajeya Cotra, a researcher at the Open Philanthropy Project and [her model now suggests](#) that it is most likely to be developed in 2040. If so, we would be able to automate the production of ideas, and indeed the production of AI systems themselves. This would allow more investment in AI, which would further increase growth, and so on. Indeed, when you plug the assumption that AI systems can substitute for all human labour into standard growth models, they predict explosive economic growth.³³

For these reasons, the prospect of explosive (>10% per year) economic growth at some point in the next 100 years is not outlandish.

It is difficult to say what the implications would be for the planet. On the one hand, high growth has been a major driver of emissions growth so far. On the other hand, the enormous technological progress would allow us to develop new technologies that could enable us to solve climate change. What exactly would happen depends on the goals of the AI systems. If they optimise for protecting the environment, then climate change would be quickly solved. But if AI systems aren't optimised to protect the environment, then emissions could increase dramatically and we might burn through all of the fossil fuels, if doing so is the cheapest way to achieve the goals of the AI system.

Thus, it would be very important to ensure that advanced AI systems optimise for socially valuable goals, or in short, that AI systems are aligned. Indeed, if one is worried about the potential climatic impacts of AI-driven explosive growth, the best way to reduce this risk

³³ Philippe Aghion, Benjamin F. Jones, and Charles I. Jones, 9. *Artificial Intelligence and Economic Growth* (University of Chicago Press, 2019); Philip Trammell and Anton Korinek, 'Economic Growth under Transformative AI' (GPI Working Paper, 2020).

would be to work on AI alignment, rather than on more traditional climate activism or advocacy. If the goals of AI systems are the main determinant of the risk of extreme emissions, then we should work to ensure that those AI systems have good goals; changing climate policy in the world more broadly will have a much smaller effect on these extreme risks. This would also help to control the many other ways, aside from climate change, that transformative AI could have extreme effects on society.

It is hard to reason about extreme and unprecedented scenarios like an AI-driven growth explosion, but I think that in that world, climate change would be orders of magnitude lower than other risks to humanity, and that the associated climate risks would be very low. None of the people who argue that AI risk is a serious problem have argued that the main risk it poses to humanity and to the future is climate change, which is some evidence for this point.

AI researchers have highlighted several different potentially concerning AI scenarios.

One scenario is advanced AI systems deliberately trying to kill humans off. But there seem to be much more efficient ways to do this than climate change, such as lethal autonomous weapons or engineered viruses.

Another scenario is that the government that had control of advanced AI systems would gain a decisive geopolitical advantage. But it seems extremely unlikely that causing extreme climate change would be the chosen way to cement international dominance because climate change would also affect the dominant country, and there seem to be much better ways to cement dominance over the international order than climate change.

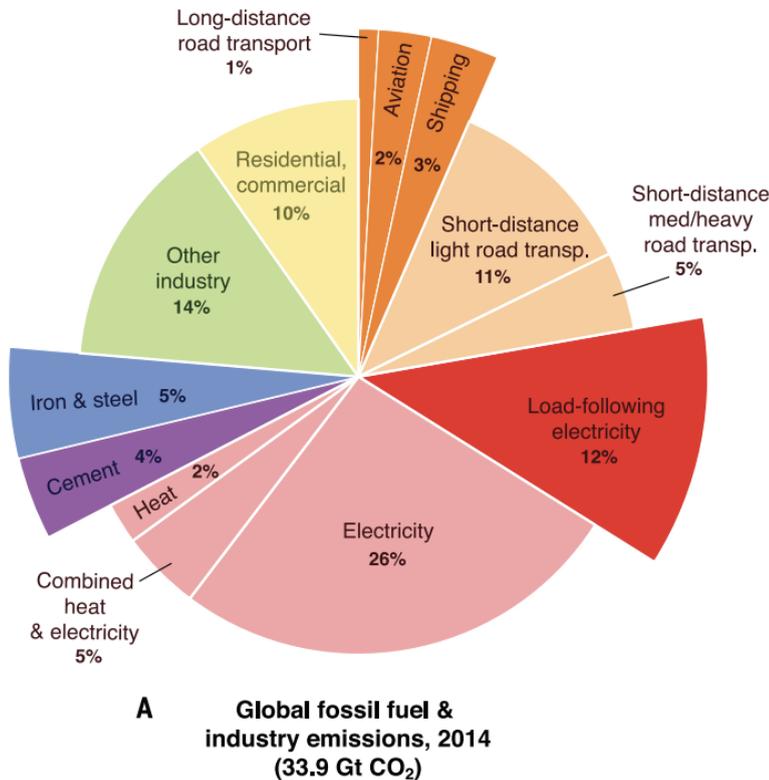
The final scenario is a world with decentralised AI systems that dramatically speed up scientific progress without killing their human controllers or giving dominance to a single government. In this world, it would have to be the case that our technological capability was far higher, but governments did not have greatly increased willingness to spend on climate change, and international coordination was not strong enough to compel countries to reduce emissions. In my view, this world seems quite unlikely to materialise. We are already seeing countries at the current technological frontier reduce emissions despite income growth and despite weak international coordination incentives. If there is dramatic technological progress, we are much more likely to be in a world in which we have a breakthrough technology that makes solving climate change much less expensive relative to today.

Even if, as I believe, there is a decent chance of an AI-driven growth explosion this century, the chance that it will lead to civilisation-destroying climate change seems slim.

1.5.3. Indefinitely stalled decarbonisation

The second scenario on which emissions might be very high is *indefinitely stalled decarbonisation*.

It now seems likely that due to progress on renewables and batteries, we will decarbonise a substantial chunk of the economy in the next few decades. But some fraction of emissions are difficult to eliminate, as shown in the chart below:



Source: Steven J. Davis et al., 'Net-Zero Emissions Energy Systems', *Science* 360, no. 6396 (29 June 2018), <https://doi.org/10.1126/science.aas9793>, Fig 2.

In brief, the reason these emissions are difficult to eliminate is that steel, cement and heavy duty transport are difficult to electrify. Steel and cement rely on high-temperature heat, which is expensive to produce with electricity. Arguably the most promising way to decarbonise these sectors is with green hydrogen, which I discuss below. Batteries are not ideal for heavy-duty transport because they are heavy and take up a lot of space.

It is hard to know what share of electricity will be provided by renewables in the future. Energy systems models suggest that costs will start to rise once the share of variable renewables in electricity passes 80%, but social and political barriers might kick in well before then, as suggested by the Cherp et al data above.

In other sectors, there is lots of scope for electrification to reduce emissions. For example, it now seems like short-distance transportation will be almost completely electrified in the next few decades, and electric heating could reduce a lot of residential and commercial emissions.

However, it remains unclear how much of the rest of the economy will be decarbonised, especially the sectors highlighted above, where decarbonisation is particularly expensive.

Energy storage

As solar and wind start to play an increasing role in energy systems, the role of energy storage is also likely to increase. The challenge of the electricity system is to ensure that demand matches supply perfectly at all times. To balance out the effects of high shares of

solar and wind, energy systems models use either controllable dispatchable sources of energy, such as gas, bioenergy, nuclear, or long-duration (multi-week or longer) energy storage.³⁴

At present, the leading form of long duration storage is pumped hydro, but this is geographically constrained. The ten largest pumped hydro storage facilities in the U.S. are collectively capable of storing a total of 43 minutes worth of U.S. energy consumption.³⁵

Batteries are not suited to providing long duration multi-week seasonal storage.³⁶ The main contenders in long-duration storage at the moment are thermal energy storage, production of hydrogen from electrolysis and storage in underground salt caverns or pressurised tanks.³⁷ There is uncertainty about how far these technologies will be scaled up, but some of them show promise.

Hydrogen

Liquid fuels such as hydrogen and ammonia (which can be made from hydrogen) are likely to be hugely important for decarbonisation. Hydrogen can provide the high-temperature heat needed to make steel and cement; ammonia can be used as a drop-in fuel in existing planes, ships and trucks; and hydrogen can be used for long-term energy storage.³⁸

Zero carbon hydrogen and ammonia are both more expensive than fossil fuel alternatives at present, but with technological progress, ammonia and hydrogen could be competitive with gasoline in the future.³⁹ To make cheap zero carbon hydrogen, we need cheap electricity and cheap catalysers. Cheap electricity could come from super-cheap renewables, but as I argued above, it is an open question whether cost declines and capacity scale-up will continue for those technologies. Another intriguing option would be to build nuclear fission hydrogen gigafactories that run 24/7 and so don't have to follow load or adjust to variable renewables, which would reduce costs.⁴⁰

Another key step would be to reduce the cost of catalysers, which can produce hydrogen through electrolysis. The costs of electrolyzers are following Wright's law: they are declining

³⁴ "Other studies partially or fully replace firm generation with one or more energy storage media capable of sustained output over weeks or longer and suited to low annual utilization rates". Jesse D. Jenkins, Max Luke, and Samuel Thernstrom, 'Getting to Zero Carbon Emissions in the Electric Power Sector', *Joule* 2, no. 12 (2018): 2498–2510.

³⁵ Jesse Jenkins and Samuel Thernstrom, "Deep Decarbonization of the Electric Power Sector: Insights from Recent Literature" (Energy Innovation Reform Project, March 2017), 6.

³⁶ Jenkins and Thernstrom, 5–6; Jenkins, Luke, and Thernstrom, "Getting to Zero Carbon Emissions in the Electric Power Sector," 10.

³⁷ "Scenarios that eschew firm generation therefore must rely upon one or more long-term energy storage technologies with an order-of-magnitude lower cost per kWh, including thermal energy storage, production of hydrogen from electrolysis and storage in underground salt caverns or pressurized tanks, or conversion of electrolytic hydrogen to methane. Considerable uncertainty remains about the real-world cost, timing, and scalability of these storage options" Jenkins, Luke, and Thernstrom, "Getting to Zero Carbon Emissions in the Electric Power Sector," 2508.

³⁸ Clean Air Task Force, "Fuels Without Carbon: Prospects and the Pathway Forward for Zero-Carbon Hydrogen and Ammonia Fuels," December 2018, <https://www.catf.us/resource/fuels-without-carbon/>.

³⁹ Clean Air Task Force, 9.

⁴⁰ Lucid Catalyst and TerraPraxis, 'Missing Link to a Livable Climate', 2020.

exponentially with additional cumulative capacity. If the trend in scale-up continues, Way et al (2021) project that costs will decline by a factor of 10 by 2050.⁴¹

Nuclear fission

Nuclear power has been behind all of the most rapid electricity decarbonisation efforts in history.⁴² If the world were to copy the per-person rate at which Sweden deployed nuclear in the 1970s and 1980s, we could decarbonise the global electricity system in less than ten years.⁴³

However, there is significant public opposition to nuclear fission in many countries due to concerns about safety, waste and weapons proliferation. These problems are greatly exaggerated. First, consider safety. Nuclear fission is actually one of the safest forms of energy production, and is [about as safe](#) as solar and wind per unit of energy produced.

Secondly, the volume of hazardous waste produced by nuclear power is small compared to other forms of waste we routinely manage.

“Whereas the ash from ten coal-fired power stations would have a mass of four million tons per year (having a volume of roughly 40 litres per person per year [in Britain]), the nuclear waste from Britain’s ten nuclear power stations has a volume of just 0.84 litres per person per year – think of that as a bottle of wine per person per year.”⁴⁴

Even in a nuclear power-reliant country, the typical per person per year volume of nuclear waste would be three orders of magnitude lower than the per person per year volume of other non-nuclear hazardous waste.⁴⁵

I discussed the link between nuclear power and nuclear weapons in [Appendix 6](#) of the 2018 Founders Pledge Climate Change Report. I argued that civilian nuclear power has historically been at most a weak driver of weapons proliferation.

In spite of this, due to its unpopularity, nuclear fission is heavily regulated, which has brought deployment to a standstill almost everywhere outside China.⁴⁶ For its 2060 decarbonisation plan, China plans to [quintuple](#) its domestic nuclear capacity. There is some chance that

⁴¹ Way et al, ‘[Empirically grounded technology forecasts and the energy transition](#)’, Oxford Martin School, 2021, p. 6.

⁴² Junji Cao et al., ‘China-U.S. Cooperation to Advance Nuclear Power’, *Science* 353, no. 6299 (5 August 2016): 547–48, <https://doi.org/10.1126/science.aaf7131>; Staffan A. Qvist and Barry W. Brook, ‘Potential for Worldwide Displacement of Fossil-Fuel Electricity by Nuclear Energy in Three Decades Based on Extrapolation of Regional Deployment Data’, *PLOS ONE* 10, no. 5 (13 May 2015): e0124074, <https://doi.org/10.1371/journal.pone.0124074>.

⁴³ “Analysis of these historical deployments show that if the world built nuclear power at no more than the per capita rate of these exemplar nations during their national expansion, then coal- and gas-fired electricity could be replaced worldwide in less than a decade.” Staffan A. Qvist and Barry W. Brook, ‘Potential for Worldwide Displacement of Fossil-Fuel Electricity by Nuclear Energy in Three Decades Based on Extrapolation of Regional Deployment Data’, *PLOS ONE* 10, no. 5 (13 May 2015): e0124074, <https://doi.org/10.1371/journal.pone.0124074>.

⁴⁴ David MacKay, *Sustainable Energy - without the Hot Air*, 169.

⁴⁵ David MacKay, 170.

⁴⁶ World Nuclear Association, *World Nuclear Performance Report 2020*, sec. 1.4.

nuclear fission could scale up elsewhere in the future as there is a good chance that many countries will struggle to meet their climate targets without it.

Nuclear fusion

In nuclear fusion, light atoms are transformed into heavier atoms to release energy, the same process that occurs within the plasma core of the Sun. Fusion produces no waste and fusion plants cannot melt down.

It is commonly joked that commercial fusion is thirty years away. Decades of public funding had seen steadily improving fusion performance, which slowed down in the 1990s.⁴⁷ However, \$2 billion of private money has recently been invested into private fusion companies, and some companies claim they will have commercially viable plants [before 2030 or 2040](#). I have no idea whether these claims are plausible or not, but there is more optimism in the field than there has been for a while.

Fossil fuels with carbon capture and storage

Carbon capture and storage involves capturing CO₂ at point sources, such as industrial chimneys, piping the CO₂ away and then injecting it underground in natural rock formations. Carbon capture and storage can be used to decarbonise:

- Electricity sector emissions from coal and gas
- Emissions from the production of high temperature heat in industrial processes
- Direct process emissions from industrial processes, such as cement.

In all of these sectors, carbon capture and storage is pure cost addition compared to not trying to capture CO₂ at all.

According to theoretical estimates, the cost to avert a tonne of CO₂ with CCS for any purpose is typically at the [very least \\$20](#) per tonne, ranging up to beyond \$100 per tonne.⁴⁸ However, these are not real world estimates, so the real world cost of CCS may be higher. Some environmentally-motivated countries have been willing to pay implicit carbon prices of upwards of \$100 per tonne, so CCS might be scaled up in some countries in the future.

In my view, the most important drawback of CCS centres on political economy. Policies that encourage CCS will probably involve giving subsidies to the fossil industry to store and sequester carbon they produce. This would keep the fossil fuel industry alive as a lobbying body with strong incentives to game the system and to oppose restrictions on carbon pollution. The real world promise of CCS very much remains to be seen.

Enhanced geothermal

[Geothermal](#) energy today provides only a tiny fraction of global energy. Conventional geothermal removes heat from shallow, water saturated rock near Earth's surface. Consequently, commercial geothermal is today confined to a few parts of the world with high-temperature heat close to the ground, such as Iceland and Japan.

⁴⁷ Anthony J. Webster, 'Fusion: Power for the Future', *Physics Education* 38, no. 2 (March 2003): 135–42, <https://doi.org/10.1088/0031-9120/38/2/305>.

⁴⁸ Global CCS Institute, "Global Cost of Carbon Capture and Storage," June 2017, 4, <https://www.globalccsinstitute.com/archive/hub/publications/201688/global-ccs-cost-updatev4.pdf>.

However, another proposed form of geothermal, known as “Super Hot Rock geothermal” has vast potential to produce low-carbon energy in much of the world. Super Hot Rock geothermal involves drilling much deeper into hot, dry crystalline rocks and then injecting water or CO₂ into these formations where high temperatures and pressure creates “supercritical” fluid that is returned to the surface for energy production. Super Hot Rock geothermal is a potentially geographically ubiquitous and cheap source of dispatchable power. Super Hot Rock geothermal remains unproven and technical barriers remain, but [some experts argue](#) that, with R&D support, the technology could become commercially viable in the next few decades. David Roberts, a climate journalist, has [argued](#) that enhanced geothermal is “an engineering problem that, when solved, solves energy”.

NET Power

The company [NET Power](#) is developing a gas power plant that captures carbon at zero additional cost, using a process called the Allam Cycle. If this works, it would be game-changing, but I am not well-placed to judge whether it will succeed.

Bioenergy

Burning sustainable bioenergy, such as sustainably managed wood or corn ethanol, is carbon-neutral in the following sense. If you burn the bioenergy, then carbon is released into the atmosphere. As the trees from the area you previously harvested grow back, they sequester the carbon again, making the process carbon neutral. Thus, although the process is not carbon neutral at the point of combustion, it becomes carbon neutral once the trees grow back.

The main downside of bioenergy is that it has large land use requirements. Current global energy demand is about 154,000 TWh. To supply a third of this with bioenergy, as is projected in some Integrated Assessment Models - would require 390 million hectares of land to be used solely for the purpose of bioenergy⁴⁹ - a third of the total arable land on the planet⁵⁰ or around 40% of the land area of the United States. This land would compete with commercial, domestic and agricultural uses. This is likely a significant barrier to the adequate scale-up of bioenergy.

Summary of technology prospects

Overall, there are a range of low carbon technologies that seem to show significant promise. However, there is some chance that technological progress in these sectors will get stuck. There are several potential reasons for this.

Excessive regulation

One possibility is that excessive regulation will kill off emerging technologies. The history of nuclear fission sets a chastening precedent. Several countries rapidly nearly completely decarbonised their electricity systems at low cost using 1970s nuclear technology. It would have been possible for the whole world, or at least the major emitters, to decarbonise

⁴⁹ Alexandre Strapasson et al., “On the Global Limits of Bioenergy and Land Use for Climate Change Mitigation,” *GCB Bioenergy* 9, no. 12 (2017): fig. 3, <https://doi.org/10.1111/gcbb.12456>.

⁵⁰ Hannah Ritchie and Max Roser, “Land Use,” *Our World in Data*, November 13, 2013, <https://ourworldindata.org/land-use>.

electricity in the coming decades by following these countries' lead. Decarbonising the rest of the economy would then have been much more straightforward.

In reality, this is not what happened. Due to some highly salient nuclear accidents, public opposition to nuclear power became intense. Nuclear regulation increased, leading to dramatically rising costs and slow construction times, particularly in Europe and the US.⁵¹ This is despite the fact that nuclear power, notwithstanding the Chernobyl disaster, is about as safe as solar and wind per unit of energy produced. The extent of nuclear regulations are discussed at length in a recent book by Jack Devanney,⁵² which is summarised in [this LessWrong post](#). It is clearly possible to build nuclear plants quickly and cheaply: in recent years, most new nuclear power plants have been built in China at low cost and typically in around 6 years.⁵³

Nuclear fission is unusually likely to be over-regulated. Nuclear accidents kill very few people on average, but they are newsworthy and highly salient to the general public. 15,000 Japanese people died in the Tohoku earthquake but most of the media attention focused on the meltdown at the Fukushima nuclear plant, which, at the time of writing, has caused [1 death from radiation and 573 from the evacuation](#).

Other emerging technologies are probably less prone to over-regulation than nuclear power. Nuclear fusion cannot meltdown and does not produce waste. As I have argued above, nuclear meltdowns and waste are trivial public health problems, but they are very important for public acceptability. It is difficult to see how energy storage could arouse much public opposition. Enhanced geothermal is renewable, cannot melt down and does not produce waste. But there might be some fracking-type opposition to this technology. NET Power might also be opposed by the environmental movement because it would allow us to continue burning fossil fuels. Solar and wind are generally popular, but also face many public acceptability barriers: the main barriers seem likely to be regulatory rather than economic or technological.

Low hanging fruit

Technological progress might be disappointing for other reasons. One concern is that ideas are getting harder to find. Exponential growth in researcher-time has been accompanied by merely constant growth in income per head. A good example is Moore's Law. The number of researchers required today to achieve the famous doubling of computer chip density is more than 18 times larger than the number required in the early 1970s. Bloom et al (2020) consider a range of sectors and find a similar picture of declining per person research productivity.⁵⁴

⁵¹ Jessica R. Lovering, Arthur Yip, and Ted Nordhaus, "Historical Construction Costs of Global Nuclear Power Reactors," *Energy Policy* 91 (April 1, 2016): 379, <https://doi.org/10.1016/j.enpol.2016.01.011>.

⁵² Jack Devanney, *Why Nuclear Power Has Been a Flop: At Solving the Gordian Knot of Electricity Poverty and Global Warming* (BookBaby, 2020).

⁵³ World Nuclear Association, *World Nuclear Performance Report 2020*, sec. 1.4.

⁵⁴ Nicholas Bloom et al., 'Are Ideas Getting Harder to Find?', *American Economic Review* 110, no. 4 (April 2020): 1104–44, <https://doi.org/10.1257/aer.20180338>.

Another related issue is declining population.⁵⁵ On endogenous growth models, economic growth and technological progress is driven by ideas. Fertility rates in most places outside of Africa are now below replacement, which means that natural population growth (not including immigration and emigration) will soon be negative in most places. China's population may already be declining. Since ideas come out of people's heads, declining fertility rates across the world could lead to long-term stagnation, especially if it proves to be harder than expected to create AI systems that can substitute for human workers.

However, in my view, it is plausible that we will have the technologies required to solve climate change in the next 30 years, long before technological progress completely stops. If we do not develop such technologies, long-run stagnation would start to become a concern and could lumber us with centuries of emissions.

Countries might also respond by implementing policies to increase the fertility rate. Rich countries could also increase immigration, which would be one way to massively increase global research productivity.

Breakdown in international coordination

It is possible that there could be a breakdown in international coordination, which causes countries to give up on their climate plans and instead to use whatever energy sources are cheapest. The most likely way this could happen is if there is a conflict between the great powers. A [range of sources](#) suggest that the chance of a Great Power conflict before 2100 is around 1 in 3.

Whether this would lead to indefinitely stalled decarbonisation depends on when the conflict would occur. If global climate ambition continues on its current trajectory, it seems likely that a substantial fraction of the problem will be solved by around 2050 anyway. Even if there is a Great Power conflict, if there are cheaper low carbon alternatives, there is little reason that we would return to fossil fuels.

Overall assessment

The best guess model used in Pielke Jnr et al (2022) finds that we would decarbonise by 2090.⁵⁶ Since this relies on IEA projections of the costs of renewables and batteries, I think this is likely pessimistic. Their models are also likely to neglect other potential breakthrough technologies, such as solar perovskites, enhanced geothermal or nuclear fusion. My best guess is that we will decarbonise completely between 2050 and 2080.

If decarbonisation does stall indefinitely, it seems unlikely that we would burn through all of the fossil fuels. As I discussed in section 1.5.1, to burn through all of the remaining fossil fuels, including the hardest to reach coal, we would have to extensive use of advanced coal extraction technology, mainly underground coal gasification, which has failed the test of commercialisation for a century and is not currently part of the energy discussion. In the

⁵⁵ Jones, C. I. (2020). The end of economic growth? Unintended consequences of a declining population (No. w26651). National Bureau of Economic Research.

⁵⁶ Roger Pielke Jr, Matthew G. Burgess, and Justin Ritchie, 'Plausible 2005–2050 Emissions Scenarios Project between 2°C and 3°C of Warming by 2100', *Environmental Research Letters* 17, no. 2 (February 2022): Fig. 4, <https://doi.org/10.1088/1748-9326/ac4ebf>.

world in which decarbonisation stalls, in order to extract all of the recoverable fossil fuels, there would have to be dramatic progress in coal extraction technology, but not in low carbon technologies which are already getting significant policy and industry attention. This is not plausible.

One possibility is that we make enough technological progress in order to develop many low carbon technologies as well as underground coal gasification this century, but progress on low carbon technology stalls just before we are able to decarbonise: perhaps decarbonising industry and heavy duty transport turns out to be too hard.

Even in this world, I find it hard to see how we could extract all of the fossil fuels. Fossil fuel extraction technology usually progresses over time, but costs are stable over time as we extract the easiest-to-reach deposits. If technological progress did truly stagnate, then we would probably stop extracting fossil fuels well before we had exhausted all recoverable resources.

One way to look at this is that so far, economic growth has been a major driver of emissions growth. Countries that burn large amounts of coal are invariably experiencing high economic growth. If economic growth stops, then we should also expect emissions growth to slow dramatically.

All of these arguments are conditioned on the assumption that the technology required to extract all of the fossil fuels would actually be difficult to create, which is what the current evidence suggests but may not turn out to be true.

In light of this argument and the arguments about explosive economic growth, it is difficult to come up with scenarios in which we would burn literally all of the recoverable fossil fuels.

If we do not use advanced coal extraction technologies, then recoverable fossil fuels are around 1,200 GtC.⁵⁷ This is below the RCP7 emissions scenario, which according to Liu and Raftery (2021) has around a 1-5% chance by 2100.

In this [simple model](#), I estimate that the chance of indefinitely stalled decarbonisation leading us to burn all of the fossil fuels is about 1 in 500,000. Perhaps the most disputable assumption in the model is the input that there is a 1 in 1,000 chance of burning all the fossil fuels conditional on technological stagnation, which one reviewer thought was too low and thought 1 in 100 was more plausible.

I also roughly calculate the time taken to burn through fossil fuels. I assume that by 2100, emissions will fall to half their current levels. If so, it would take around 400 years to burn through all of the fossil fuels.

⁵⁷ This is using Ritchie and Dowlatabadi's estimates. For calculations, see the [sheet](#) on ultimately recoverable fossil fuels.

1.6. Can we change course?

One counter to the argument that climate change poses serious global catastrophic risks is that we can learn as we go. If climate sensitivity turns out to be higher than expected, or impacts turn out to be worse than expected, then we can change course and so avoid the very worst outcomes. How easy would it be to change course?

1.6.1. The physics

It is important to first clarify some physics. Reducing temperatures does not just merely involve reducing emissions rates. Temperatures increase linearly with cumulative CO₂ emissions, so as long as emissions are positive, temperatures will continue to increase. In order to reduce temperatures, we need to reduce emissions to zero *and* remove CO₂ from the atmosphere.

The speed at which we can remove CO₂ from the atmosphere is limited by the fact that when we remove carbon from the air, natural sinks of carbon (the oceans and the land) release carbon back into the atmosphere. So, sucking out and storing one tonne of carbon from the atmosphere does not cause there to be one tonne less of carbon in the atmosphere, once the system adjusts. Over a period of several decades, this would replace up to half of the CO₂ that had been removed by negative emissions.⁵⁸ If we ever want to return CO₂ concentrations back to their pre-industrial level, we will eventually have to suck out all of the carbon we have ever emitted.

Emitting CO₂ is like adding concrete blocks to a tower.⁵⁹ Each year we emit, concentrations get higher, and the tower gets taller. Even if we stop adding blocks altogether, we still have to figure out what to do with the tower. Unless we dismantle the tower ourselves, it will only be eroded away after tens of thousands of years. Even after 100,000 years, remnants of the tower will remain.

1.6.2. How hard is changing course?

Emissions reduction should be the priority

The realities of climate physics mean that if we want to return CO₂ concentrations back to safe levels, we not only have to stop emitting, but also to start sucking CO₂ out of the atmosphere ourselves. In my view, it would be rational to *first* focus on reducing emissions before doing negative emissions.

Consideration of the basic physics illustrates why this is so. Although CO₂ concentrations are high in terms of recent history, CO₂ is still highly diffuse - atmospheric concentrations at the moment are around 415 parts per million, or 0.04%. In contrast, the chimneys of fossil fuel plants often contain gas that is 10% CO₂.⁶⁰ concentrations are 250x higher than in the

⁵⁸ National Academy of Sciences, *Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration* (Washington, D.C.: National Academies Press, 2015), 29–32, <http://www.nap.edu/catalog/18805>.

⁵⁹ I owe this analogy to a tweet by Glen Peters.

⁶⁰ Oliver Morton, *The Planet Remade: How Geoengineering Could Change the World* (London: Granta, 2015), 245–46.

ambient air today.⁶¹ This suggests it will generally be easier to not emit CO₂ in the first place than to suck it back out of the air.

This reasoning is confirmed by cost estimates of negative emissions. The only highly scalable negative emissions technology is direct air capture, but that costs hundreds of dollars per tonne of CO₂ sequestered, and one recent review suggests that the costs of direct air capture will never fall below \$100 per tonne.⁶² In contrast, renewables and nuclear can in principle substitute for fossil fuels at relatively low cost.

Even if one country reduces their emissions to zero, it would, if possible, make more sense for them to pay other countries to stop burning fossil fuels than to start doing negative emissions. However, historically, international offset programmes have been beset by problems, and many argue that they do not bring real climate benefits.⁶³ I have looked at this in some depth and believe that most international offsets bring little environmental benefit. So, it might make sense for states or groups of states to unilaterally deploy negative emissions technology once they have reduced emissions to zero, even if other states are still emitting.

1.6.3. Crash direct air capture

There are many different forms of negative emissions technology, and all of them have technological, economic and political limitations. The only form of negative emissions technology that is well-suited to a drastic course correction is direct air capture and storage, which captures carbon directly from the ambient air and stores it underground. This is because:⁶⁴

- Deployments are modular, scalable and highly controllable by the governments and firms that invest
- Carbon removals are verifiable
- Deployment does not harm existing special interests or consumers
- Though energy-intensive, direct air capture has no other biophysical limits
- It does not require large land use change, which would compete with agricultural and other uses
- Unlike strategies for controlling emissions from industry and the broader economy, deploying direct air capture does not intrinsically require intrusive policy interventions that might damage a country's economic competitiveness.
- It can be deployed unilaterally

Once we have decarbonised the economy, getting to net negative emissions would be straightforward as we just could use spare low carbon energy for direct air capture.

⁶¹ Oliver Morton, *The Planet Remade: How Geoengineering Could Change the World* (London: Granta, 2015), 245–46.

⁶² Sabine Fuss et al., 'Negative Emissions—Part 2: Costs, Potentials and Side Effects', *Environmental Research Letters* 13, no. 6 (May 2018): Table 2, <https://doi.org/10.1088/1748-9326/aabf9f>.

⁶³ Raphael Calel et al., 'Do Carbon Offsets Offset Carbon?', 2021; Kevin Anderson, 'The Inconvenient Truth of Carbon Offsets', *Nature* 484, no. 7392 (2012): 7–7.

⁶⁴ Ryan Hanna et al., 'Emergency Deployment of Direct Air Capture as a Response to the Climate Crisis', *Nature Communications* 12, no. 1 (14 January 2021): 2, <https://doi.org/10.1038/s41467-020-20437-0>.

Hanna et al (2021) outlines what would be involved in an emergency mass deployment of direct air capture. They assume:

- **Scenarios:** Hanna et al consider three scenarios of crash direct air capture deployment.
 - By the US only
 - By the OECD group of rich democracies
 - By International Bank for Reconstruction and Development members, which includes some low- and middle-income countries.
- **Spending:** The actors spend \$1-1.6 trillion per year, which then grows with GDP over time.
 - The US unilaterally spends 5% of its GDP, comparable to average wartime deployment.
 - OECD countries spend 2% of GDP.⁶⁵
 - International Bank for Reconstruction and Development members spend \$1.6 trillion.⁶⁶ I'm not sure what this is as a percentage of IBRD GDP.
- **Technology costs**
 - Costs of direct air capture are assumed to initially be \$150 to \$1000 per tonne of CO₂, but they fall to \$75 to \$600 per tonne by 2075.⁶⁷
- **Scale-up speed**
 - Industry scale-up is limited to 20% per year, which is comparable to the recent growth in solar photovoltaics, the Liberty ship building programme in the US in World War 2, and the French nuclear programme in the 1970s.
 - The growth rate is 20% per year in 2025 to 2050, 5% 2050 to 2075, and 1% in 2075 to 2100.⁶⁸

The chart below shows the effect this has on emissions and temperatures

⁶⁵ The states spend \$1.4 trillion. OECD [GDP is \\$64 trillion](#). Ryan Hanna et al., 'Emergency Deployment of Direct Air Capture as a Response to the Climate Crisis', *Nature Communications* 12, no. 1 (14 January 2021): SI p9, <https://doi.org/10.1038/s41467-020-20437-0>.

⁶⁶ The states spend \$1.6 trillion. IBRD GDP is Hanna et al., 10.

⁶⁷ Hanna et al., 'Emergency Deployment of Direct Air Capture as a Response to the Climate Crisis', Fig. 5.

⁶⁸ Hanna et al., 'Emergency Deployment of Direct Air Capture as a Response to the Climate Crisis', Fig. 3.

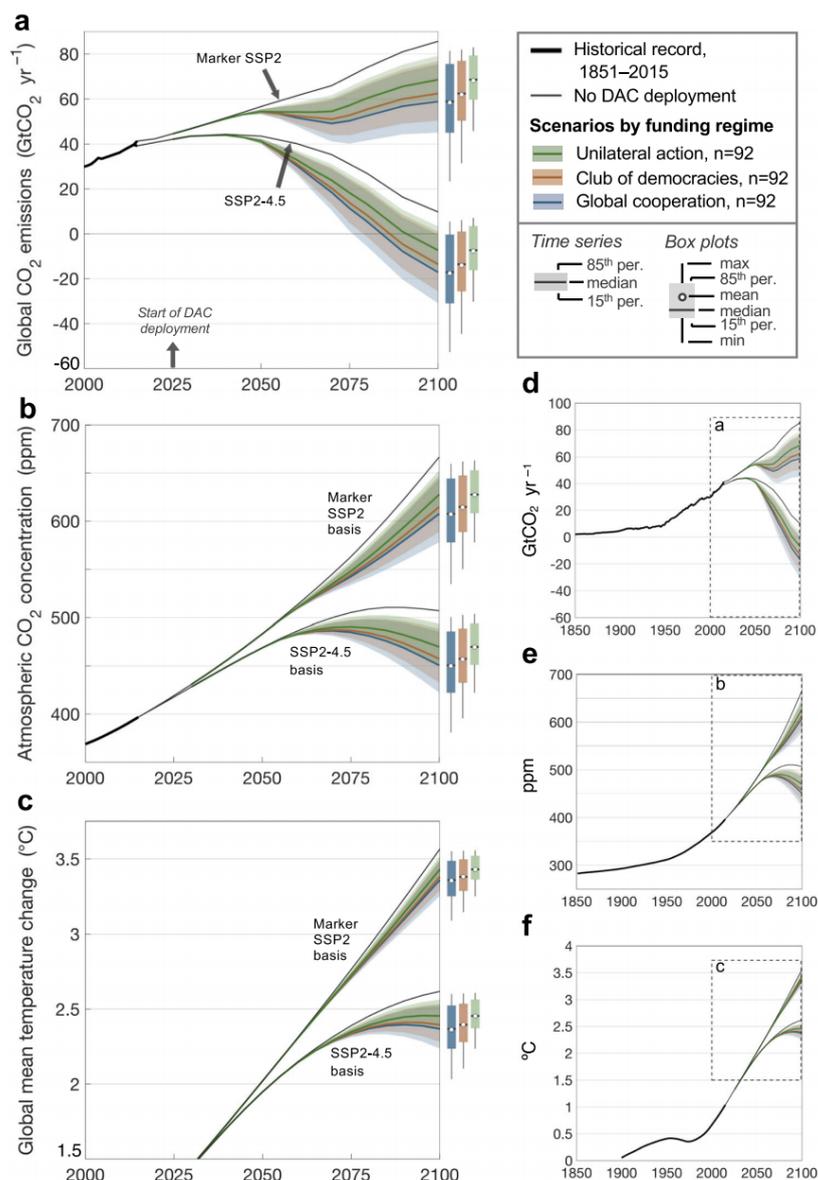


Fig. 4 Climate benefits of net CO₂ removal assuming marker SSP2 and SSP2-4.5 emission futures. **a** Global CO₂ emissions. SSP emissions are from the SSP database version 2.0 (<https://tntcat.iiasa.ac.at/SspDb>; ref. 49). **b** Atmospheric CO₂ concentration. **c** Global mean temperature change relative to pre-industrial levels (1850–1900). **d–f** Trends beginning in 1850. Ribbons indicate the 15th and 85th percentile scenarios; solid lines indicate the median. Black lines show the case of no DAC deployment. Boxes show the 15th and 85th percentile scenarios in 2100; center line, median; dot, mean; whiskers, range. (Here we show SSP2, IPCC’s middle-of-the-road pathway; we include scenarios with a wider range of emissions—SSP1-2.6 and SSP5-8.5—in Supplementary Fig. 8).

By the end of the century, the OECD Direct Air Capture programme is removing 7 billion tonnes of carbon each year, and from 2025 to 2100, it will remove 0.2 trillion cumulative tonnes of carbon.⁶⁹ By 2150, temperatures are 0.7°C lower relative to no DACs, halving the warming that would have happened without direct air capture.⁷⁰

⁶⁹ Hanna et al., ‘Emergency Deployment of Direct Air Capture as a Response to the Climate Crisis’, Fig. 3.

⁷⁰ “Nevertheless, the effect of DAC on the temperature trajectory is substantial—it arrests the growth in the warming curve, which peaks at 2.4–2.5 °C in the 2090s. For model runs that extend further into the future, DAC reverses temperature rise to 1.9–2.2 °C in 2150, a reversal of 38–61% of the warming that occurs without DAC, which sees temperatures reach 2.7 °C in 2150 and rise even further thereafter.” Hanna et al., ‘Emergency Deployment of Direct Air Capture as a Response to the Climate Crisis’.

This also allows us to work out the effects of greater levels of direct air capture on emissions reduction and on temperatures.⁷¹ Below I assume that deployment stays at 2100 levels indefinitely, which will be close to correct as Hanna et al (2021) assume that capacity growth slows substantially by 2100.

Cost as a % of the GDP of OECD democracies	Cumulative emissions reduced (trillion tonnes of carbon per century)	Warming reduced by (°C)
2%	0.3	0.5
6%	0.9	1.4
8%	1.2	1.8
20%	3	4.5

Calculations are in [this guesstimate model](#).

This illustrates that we *can* reverse climate change if we are willing to make significant investments. If OECD countries spent 6% of GDP, they would capture 900 billion tonnes of carbon, which is greater than total carbon emissions on RCP4.5.

1.6.4. How likely is crash direct air capture?

It is clear that crash direct air capture is technically feasible, but how likely is it? Deployment of negative emissions technology today is negligible.⁷² At the moment, typical costs are estimated to be more than \$250 per tonne, though they might fall to around \$100 per tonne with enough support. Would countries ever be willing to pay this much to reduce emissions?

Observed willingness to pay to abate

One natural way to assess willingness to pay to abate carbon is to look at observed carbon prices. As of 2019, the global carbon price was [around \\$2 per tonne](#). The global price now is likely somewhat higher due to increases in the EU Emissions Trading Scheme price, but I think it is still below \$10 per tonne. Indeed, once we take account of fossil fuel subsidies, the *net* global carbon price is probably *negative*. The Information Technology and Innovation

⁷¹ “The cooling (or avoided warming) due to CDR would be proportional to the cumulative amount of CO₂ removed from the atmosphere by CDR (Tokarska and Zickfeld, 2015; Zickfeld et al., 2016), as implied by the near-linear relationship between cumulative carbon emissions and GSAT change” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, Ch. 3 sec. 4.6.3.2.

⁷² “In the United Kingdom, Drax has begun a pilot project to capture CO₂ from its biomass-fuelled power plant. If the project is successful, it could become the world’s first negative emissions power plant” “Today, more than 10 direct air capture plants are operating in Europe, the United States and Canada. Most of these plants are small and sell the captured CO₂ for use – for carbonating drinks, for example. However, the first large-scale direct air capture plant is now being developed in the United States by a partnership between Carbon Engineering and Occidental Petroleum. The plant will capture up to 1 million tonnes of CO₂ each year for use in enhanced oil recovery and could become operational as early as 2023. In Iceland, the CarbFix project is capturing CO₂ from the atmosphere for injection and storage in basalt rock formations.” IEA, [Going carbon negative: what are the technology options?](#) Jan 2020

Foundation estimates that in 2019, the effective net global carbon price was *minus* \$10 per tonne.⁷³

However, global carbon prices are not a good indicator of willingness to pay to reduce emissions because carbon pricing faces especially severe political economy problems relative to other climate policies: they have highly visible costs, and are imposed on sectors with very different abatement costs and with very different levels of political organisation.

Because the political economy barriers to carbon pricing are so severe, the price in global carbon pricing schemes is not a good measure of observed willingness to pay to reduce emissions. Many policies imposed in many jurisdictions have implicit carbon prices in excess of \$100 per tonne, and even of \$1,000 per tonne.

- The decision to close still viable coal plants in Ontario had an implicit carbon price of \$80-\$100 per tonne.⁷⁴
- The implicit carbon price of Germany's solar power incentives was upwards of \$500 per tonne.⁷⁵
- The implicit carbon price of a Norwegian suite of regulations on petrol cars amounts to an implicit carbon price of more than 1000 Euros per tonne.⁷⁶

So, it is not implausible that some rich environmentally-motivated democracies would be willing to pay more than \$100 per tonne to remove carbon through direct air capture.

How might willingness to pay change in the future?

Willingness to pay might also increase in the future because (1) incomes will be much higher, and (2) the impacts of climate change will be worse. On the middle of the road shared socioeconomic pathway 2, GDP per capita is \$80,000 in 2100, compared to around \$10,000 today. On the high growth SSP5-8.5, GDP per capita is \$140,000 in 2100. Since environmental protection is plausibly a luxury good, we should expect people to be willing to spend a greater share of their income on it as their income rises.

Secondly, crash direct air capture would be more likely to happen if climate change were starting to have especially bad impacts on certain countries; Hanna et al (2021) explicitly envision the crash direct air capture programme in the case of a 'climate emergency'.

There are cases of individual countries spending huge fractions of GDP in times of crisis. COVID-19 is one example. During World War II, military spending was upwards of 20% of

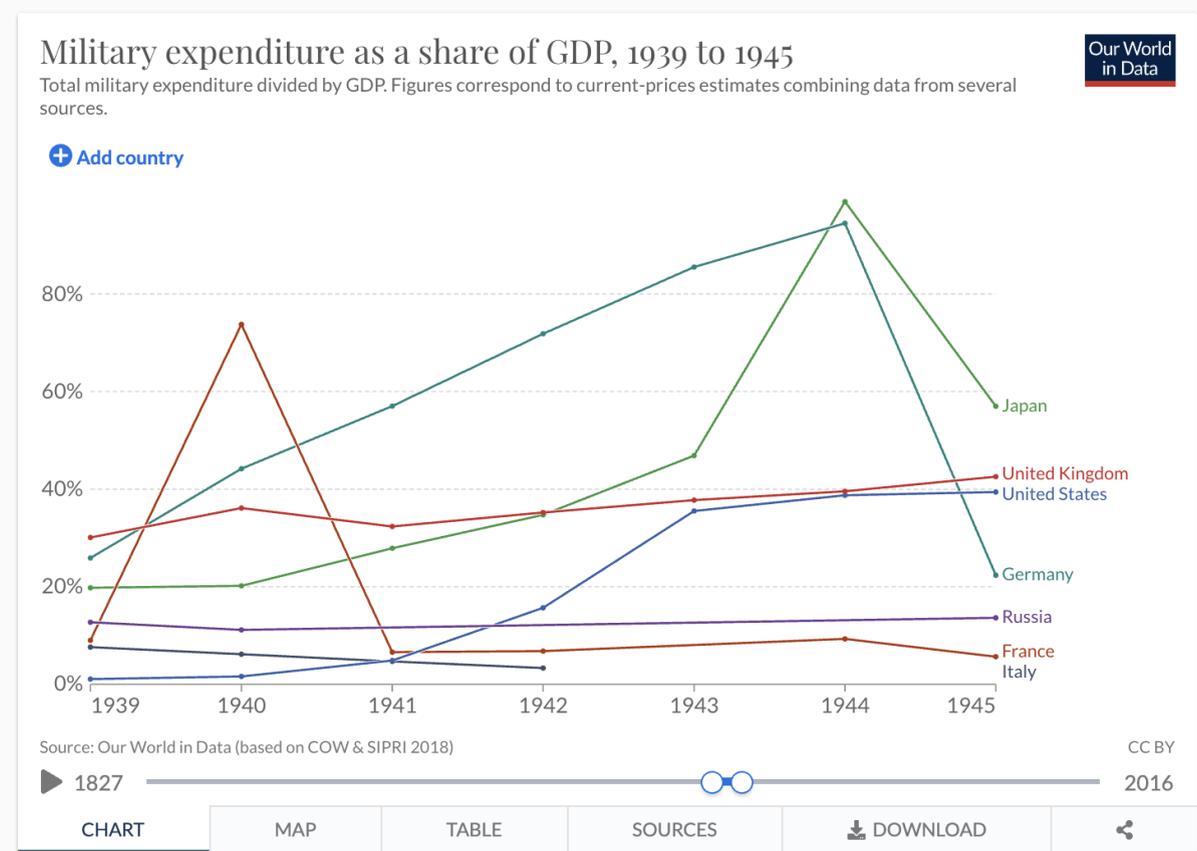
⁷³ See endnote 24 of the [ITIF report Omission Innovation 2.0 \(2019\)](#).

⁷⁴ See this [review](#) by Jaccard et al.

⁷⁵ "Results show that for the period analyzed both the RE carbon surcharge and the implicit carbon price of wind are on the order of tens of euro per tonne of CO₂, while for solar are on the order of hundreds of euro per tonne of CO₂." Claudio Marcantonini and A. Denny Ellerman, 'The Implicit Carbon Price of Renewable Energy Incentives in Germany', *The Energy Journal* 36, no. 4 (1 October 2015), <https://doi.org/10.5547/01956574.36.4.cmar>.

⁷⁶ "The price of carbon characterizing the trade-off between conventional and battery electric cars in Norway as of 2019 exceeds €1370 per ton of CO₂." Lasse Fridstrøm, 'The Norwegian Vehicle Electrification Policy and Its Implicit Price of Carbon', *Sustainability* 13, no. 3 (January 2021): 1346, <https://doi.org/10.3390/su13031346>.

GDP for most of the major powers. In Japan and Germany, it surpassed 90% of GDP in 1944.



A key question then is how severe climate change might be relative to World War II for countries with the spending power to invest in crash direct air capture. These will most likely include current OECD countries as well as many emerging economies in Asia.

I argue in the remainder of this report that warming of 4°C would pose a much smaller threat to most countries than world war. Small island nations could face an existential threat, and some of these may have high per capita GDP, but they would lack the financial resources to carry out a direct air capture programme at any scale. The literature on warming of >5°C is much thinner, so there is more uncertainty about how bad impacts would be at this level of warming.

Since World War II, American defence spending during wartime (for wars in Korea, Vietnam, the Persian Gulf, Iraq, and Afghanistan) has generally been 1-10% of GDP.⁷⁷ It is perhaps more plausible that spending on negative emissions could reach this level if warming passed 3°C, or if a tipping point appeared to have been crossed. Spending 6% of the GDP of the OECD countries over a century would reduce temperatures by around 1.5°C.

It is difficult to say exactly when states would be motivated to engage in crash direct air capture. My best guess is that efforts would start to increase once warming passed 3°C to

⁷⁷ Stephen Daggett, 'Costs of Major U.S. Wars' (Congressional Research Service, 29 June 2010), <https://sgp.fas.org/crs/natsec/RS22926.pdf>, Table 1.

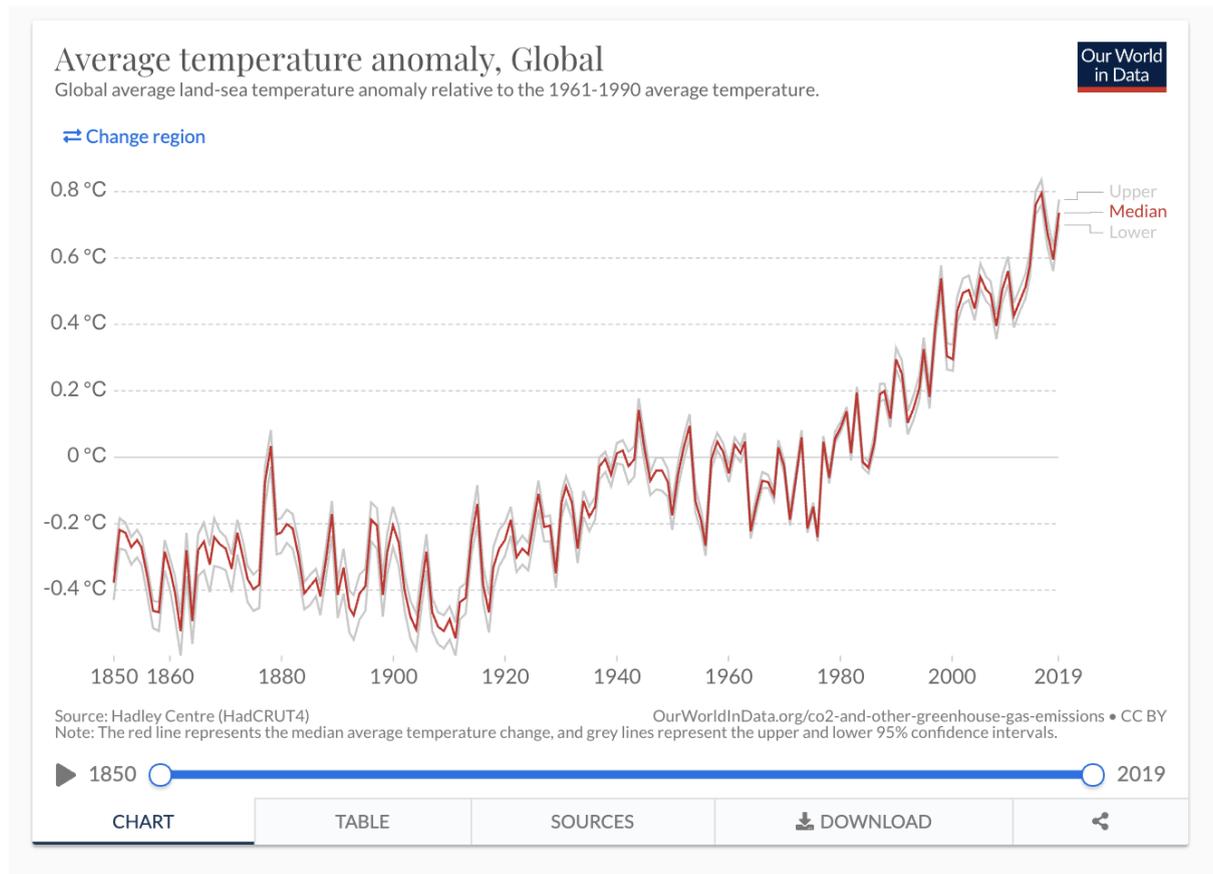
4°C. This level of warming would create very high levels of heat stress for countries in the tropics and subtropics, including regions that are likely to have a lot of economic might in the future such as the US, China, India and Europe. At 6°C of warming, New York, New Orleans and Washington DC would face higher heat stress than modern day Bahrain. At 4°C of warming, huge population centres in India and China would face heat stress only today seen occasionally in the Persian Gulf.⁷⁸

Although these countries would be able to adapt with air conditioning, it would still be unpleasant to be outside in the summer months, and would make outdoor exercise impossible. I would guess that at >4°C, the chance of decarbonisation and a direct air capture programme costing >2% of the GDP of certain major powers is upwards of 50%.

⁷⁸ On this, see Chapter 6.

2. How hot will it get?

Since the Industrial Revolution, temperatures have increased by about one degree Celsius, with most of the increase coming after 1980.



2.1. Climate sensitivity metrics

There are many different metrics that quantify the effect of greenhouse gas emissions on the climate. These can be defined as follows:⁷⁹

Equilibrium climate sensitivity = The global average warming we get following a doubling of atmospheric CO₂ once all climate processes except ice sheets reach equilibrium. This can be roughly thought of as the warming we get over several centuries, assuming that CO₂ concentrations remain at a certain level.

Earth system sensitivity = The warming we get following a doubling of atmospheric CO₂ once the system reaches equilibrium over several millennia. This accounts for ice sheet feedbacks.

Transient climate response to cumulative emissions = The warming we get following a given amount of cumulative emissions.

⁷⁹ There are also other metrics such as *effective climate sensitivity*, and *transient climate response*. S. Sherwood et al., 'An Assessment of Earth's Climate Sensitivity Using Multiple Lines of Evidence', *Reviews of Geophysics*, 2020, e2019RG000678.

In climate science, most of the attention has been focused on estimating equilibrium climate sensitivity. The equilibrium climate sensitivity metric rests on the insight that there is a **logarithmic** relationship between atmospheric CO₂ *concentrations* and warming. Each doubling of CO₂ concentrations produces the same amount of warming. Doubling CO₂ concentrations from 280ppm to 560ppm, produces approximately the same warming as doubling them again to 1,120ppm (with some caveats discussed below). Thus, there are “diminishing returns” from concentrations to warming. Earth system sensitivity rests on a similar insight but considers a longer timeframe for the earth system to reach equilibrium.

However, the Transient Climate Response to Cumulative Emissions suggests that there is a **linear** relationship between *emissions* and warming. How can both equilibrium climate sensitivity and the transient climate response to cumulative emissions be correct? How can there be a logarithmic relationship between concentrations and warming, but a linear relationship between emissions and warming?

The answer is that the diminishing effect of higher CO₂ concentrations is compensated by the diminishing ability of the ocean to take up heat and carbon.⁸⁰ The amount of CO₂ we release that stays in the atmosphere - the ‘airborne fraction’ - changes depending on past emissions and on temperatures. At present, the airborne fraction is around 45% - the rest is absorbed by the oceans and the land. However, as emissions increase, carbon sinks also become less effective at removing CO₂ from the atmosphere, which results in a higher airborne fraction. This is shown in the chart below:

⁸⁰ “The TCRE emerges from the diminishing radiative forcing from CO₂ per unit mass being compensated for by the diminishing ability of the ocean to take up heat and carbon” Thomas L. Frölicher and David J. Paynter, ‘Extending the Relationship between Global Warming and Cumulative Carbon Emissions to Multi-Millennial Timescales’, *Environmental Research Letters* 10, no. 7 (2015):

Total cumulative CO₂ emissions taken up by land and ocean (colours) and remaining in the atmosphere (grey) under the five illustrative scenarios from 1850 to 2100

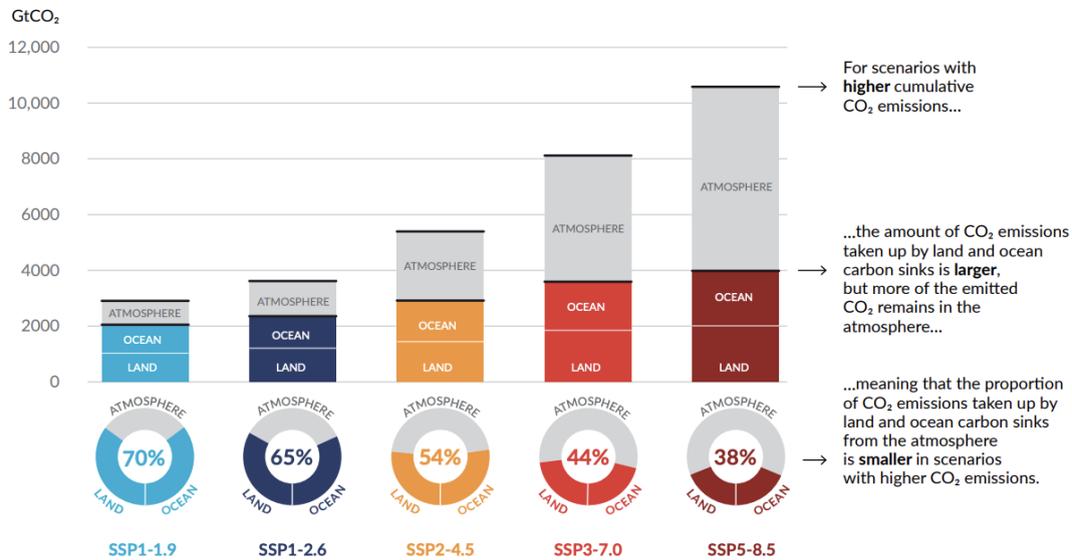


Figure SPM.7 | Cumulative anthropogenic CO₂ emissions taken up by land and ocean sinks by 2100 under the five illustrative scenarios

The cumulative anthropogenic (human-caused) carbon dioxide (CO₂) emissions taken up by the land and ocean sinks under the five illustrative scenarios (SSP1-1.9, SSP1-2.6, SSP2-4.5, SSP3-7.0 and SSP5-8.5) are simulated from 1850 to 2100 by Coupled Model Intercomparison Project Phase 6 (CMIP6) climate models in the concentration-driven simulations. Land and ocean carbon sinks respond to past, current and future emissions; therefore, cumulative sinks from 1850 to 2100 are presented here. During the historical period (1850–2019) the observed land and ocean sink took up 1430 GtCO₂ (59% of the emissions).

Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, SPM. 8.

In addition, the oceans take up less heat as they warm. So, in a nutshell, there are:

- Increasing returns from emissions to concentrations
- Decreasing returns from concentrations to warming

These effects cancel out the logarithmic effect of atmospheric carbon on warming, such that warming is found to be proportionate to cumulative emissions in Earth System Models as shown in the chart below:

Every tonne of CO₂ emissions adds to global warming

Global surface temperature increase since 1850–1900 (°C) as a function of cumulative CO₂ emissions (GtCO₂)

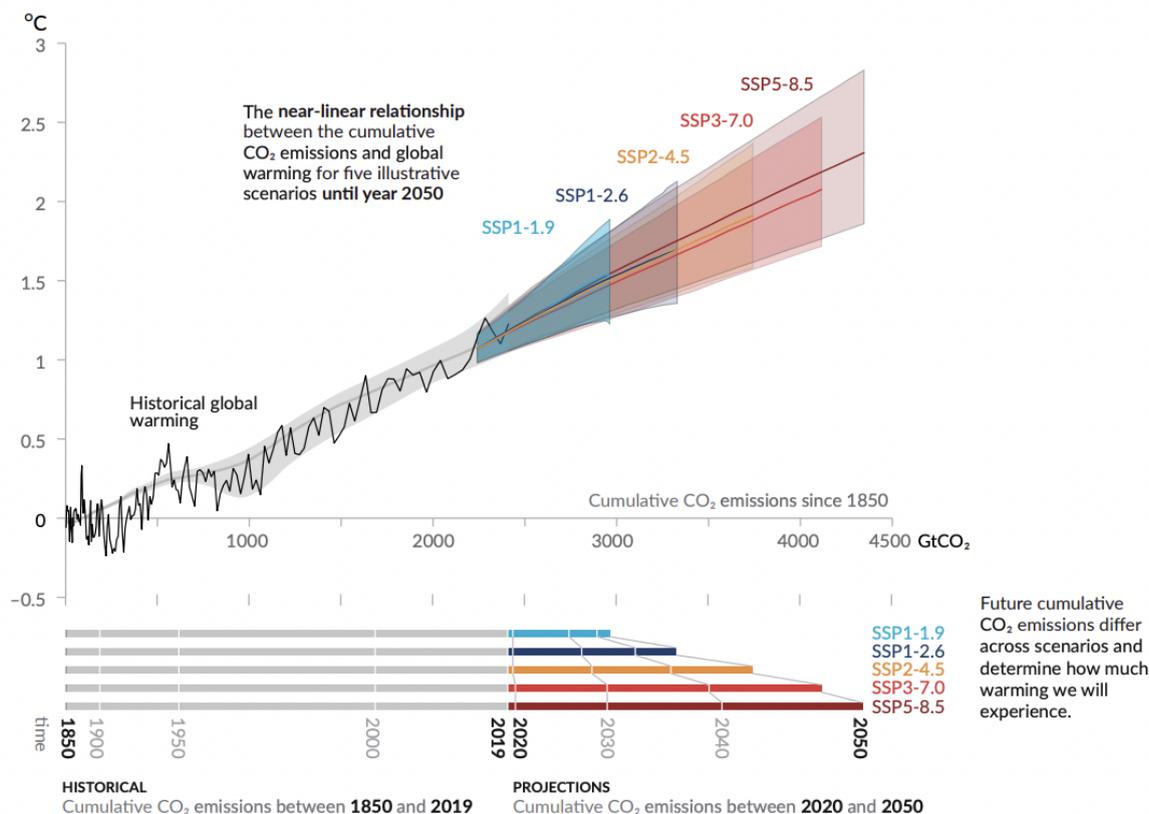


Figure SPM.10 | Near-linear relationship between cumulative CO₂ emissions and the increase in global surface temperature

Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, Fig. SPM. 10.

2.1.1. Advantages and disadvantages of different metrics

Each of these metrics has advantages and disadvantages.

Equilibrium metrics

In spite of its prominence, equilibrium climate sensitivity is only indirectly practically useful. It measures the warming we get on the assumption that CO₂ concentrations reach a certain level and then stay there indefinitely. But, once emissions stop, CO₂ concentrations would actually decline due to natural uptake of CO₂, initially by the oceans. Atmospheric CO₂ over time would look more like this after a given level of emissions over one million years:

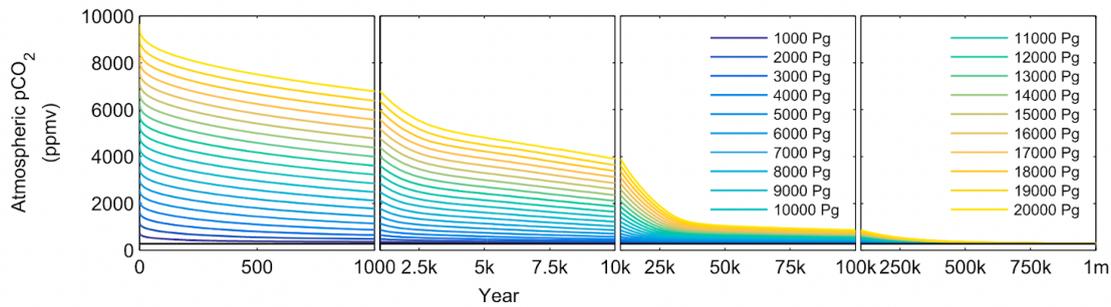


Figure 2. Atmospheric $p\text{CO}_2$ predicted by cGENIE for the pulse series scenarios (1000–20,000 Pg C). Preindustrial CO_2 concentrations are shown in black.

Source: N. S. Lord et al., ‘An Impulse Response Function for the “Long Tail” of Excess Atmospheric CO_2 in an Earth System Model’, *Global Biogeochemical Cycles* 30, no. 1 (2016): 2–17, <https://doi.org/10.1002/2014GB005074>.

CO_2 concentrations would only stay at a certain level indefinitely if there were a very precise low level of CO_2 emissions sustained over many centuries to precisely compensate for ocean CO_2 uptake.⁸¹ This is unlikely to happen in the real world. It is more plausible that emissions would cease altogether, or that emissions would decline over a long period by a constant percentage. So, for plausible real world emissions scenarios, warming is not related to equilibrium metrics in a simple way.

Earth system sensitivity is a more complete metric than equilibrium climate sensitivity because it includes ice sheet loss. It takes thousands of years for ice sheets to fully melt following warming, and because water and the land uncovered by ice melt are darker than ice, planetary reflectivity is reduced, which traps more heat. For this reason, earth system sensitivity is thought to be 1-2 times greater than equilibrium climate sensitivity.⁸² Like equilibrium climate sensitivity, earth system sensitivity assumes that concentrations remain at a precise point indefinitely, which is unlikely to happen in the real world.

Transient Climate Response to Cumulative Emissions

The advantage of the transient climate response to cumulative emissions metric is that it includes changes in the carbon cycle, so it does not require emissions to be tweaked to maintain a constant level of concentrations. Rather, knowing the transient climate response

⁸¹ “Physically this can be understood by realizing that the ECS is a theoretical quantity representing the warming that would occur only if atmospheric concentrations of greenhouse gases were held constant indefinitely while the climate system was allowed to come into equilibrium. Such a ‘constant radiative forcing’ scenario would require a very precise low level of emission of CO_2 sustained over many centuries to precisely compensate for ocean CO_2 uptake. This is clearly not a particularly policy-relevant scenario.” Richard Millar et al., ‘The Cumulative Carbon Budget and Its Implications’, *Oxford Review of Economic Policy* 32, no. 2 (2016): 323–42.

⁸² “Since we do not have accurate estimates of the ice sheet and vegetation forcings, we instead use an uncertain parameter to represent the amount by which these (generally slower) responses inflate the response that would be generated by CO_2 alone. Lunt et al. (2010) argue that this ratio ESS/S is around 1.4 for the Pliocene based on simulations using HadCM3, while Haywood et al. (2013) find an ensemble mean ratio of 1.5 with considerable variation between models but with a total range of 1 to 2 across the models in the PlioMIP1 ensemble. We represent these results with an ESS inflation factor $1 + f_{\text{ESS}}$ where f_{ESS} is distributed as $N(0.5, 0.25)$.” S. Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’, *Reviews of Geophysics*, 2020, p. 59.

to cumulative emissions, we can infer from cumulative emissions how much warming there is going to be. Because models of the TCRE include the carbon cycle, they can account for how CO₂ is drawn down by natural processes over time, whereas equilibrium metrics necessarily hold CO₂ concentrations constant.

Once emissions stop, temperature would stay roughly constant for about 100 years⁸³ before slowly declining over hundreds of thousands of years.

One downside of the TCRE is that it is not clear whether the relationship holds for emissions beyond 1 trillion tonnes of carbon.⁸⁴ Another drawback is that estimates of transient climate response to cumulative emissions do not usually include climate system feedbacks that are relevant on multi-millennial timescales, such as ice sheets and vegetation.

Overall, TCRE is a more useful metric than equilibrium climate sensitivity. It is easier to estimate warming using the Transient Climate Response to Cumulative Emissions, and the real world implications of equilibrium climate sensitivity are unclear.

2.2. Estimates of climate sensitivity

2.2.1. Equilibrium Climate Sensitivity

Of all climate sensitivity metrics, equilibrium climate sensitivity has been the main focus of scientific attention so far. The Swedish scientist Svante Arrhenius was the first person to quantify equilibrium climate sensitivity. In 1896 he estimated that for each doubling of CO₂ concentrations, the world would warm by between 5°C and 6°C. In 1906, he revised this down to 4°C, which is remarkably close to modern estimates of climate sensitivity.⁸⁵

⁸³ “The Zero Emissions Commitment (ZEC) is the change in global mean temperature expected to occur following the cessation of net CO₂ emissions and as such is a critical parameter for calculating the remaining carbon budget... Overall, the most likely value of ZEC on multi-decadal timescales is close to zero, consistent with previous model experiments and simple theory.” Andrew H. MacDougall et al., ‘Is There Warming in the Pipeline? A Multi-Model Analysis of the Zero Emissions Commitment from CO₂’, *Biogeosciences* 17, no. 11 (15 June 2020): Figure 3b.

<https://doi.org/10.5194/bg-17-2987-2020>.

⁸⁴ “Overall, there is high agreement between multiple lines of evidence (robust evidence) resulting in high confidence that TCRE remains constant for the domain of increasing cumulative CO₂ emissions until at least 1500 PgC, with medium confidence of it remaining constant up to 3000 PgC because of less agreement across available lines of evidence.” IPCC, *Climate Change 2021: The Physical Science Basis*, Ch. 5, sec. 5.5.1.2.1.

⁸⁵ “The citations of Arrhenius’ calculations are usually based on the work published in 1896 [Arrhenius, 1896]. In this paper Arrhenius reported that CO₂ doubling should increase the Earth’s mean temperature by 5-6°C. In the same year, he estimated that it would take about 3000 years for mankind to double the atmospheric concentration through the burning of fossil fuels [Arrhenius, 1896]. However, the 1896 paper was not his last publication on the problem of global warming. In later works, Arrhenius revised the estimates mentioned above. It is not clear exactly how he derived his values, but these later values are much closer to modern estimates than most think. For example, in the 1906 book, *Worlds in the Making: The Evolution of the Universe* [Arrhenius, 1906], Arrhenius wrote that “...any doubling of the percentage of carbon dioxide in the air would raise the temperature of the Earth’s surface by 4°C...”” Andrei G. Lapenis, ‘Arrhenius and the Intergovernmental Panel on Climate Change’, *Eos, Transactions American Geophysical Union* 79, no. 23 (1998): 271–271.

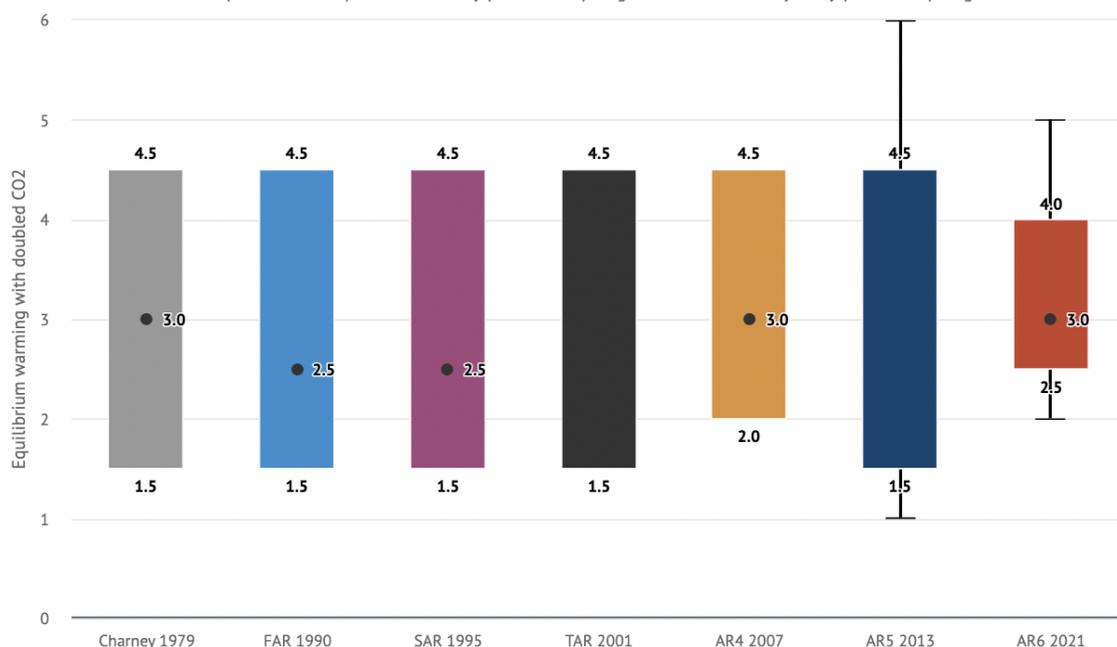
For many decades after the 1970s, estimates of equilibrium climate sensitivity barely changed. In the IPCC's 2013-14 Fifth Assessment Report, the 95th percentile of climate sensitivity was around 6°C.⁸⁶

This was a major motivation for arguments by the late great economist Martin Weitzman that climate change is a serious global catastrophic risk. In their 2015 book *Climate Shock*, Wagner and Weitzman argued that, based on what was then thought to be current policy (RCP6), there is an 11% chance of more than 6°C of warming.⁸⁷

However, the IPCC's Sixth Assessment Report has since narrowed the 90% confidence interval for climate sensitivity at the lower and higher end, as shown in the diagram below

AR6 narrows the range for equilibrium climate sensitivity and 'rules out' low values

Dots show central estimates (when available). Bars show likely (66% chance) range. Whiskers show very likely (5% to 90%) range.



Source: CarbonBrief, IPCC: How the AR6 WG1 summary for policymakers compares to its predecessor

The main reason that uncertainty about climate sensitivity has narrowed is because it is now calculated using multiple lines of evidence.⁸⁸ This is done in a formal Bayesian way in

⁸⁶ Though the IPCC wasn't completely clear about this.

⁸⁷ Gernot Wagner and Martin L. Weitzman, *Climate Shock: The Economic Consequences of a Hotter Planet* (Princeton: Princeton University Press, 2015), Table 3.1.

⁸⁸ "In AR6, the assessments of ECS and TCR are made based on multiple lines of evidence, with ESMS representing only one of several sources of information. The constraints on these climate metrics are based on radiative forcing and climate feedbacks assessed from process understanding (Section 7.5.1), climate change and variability seen within the instrumental record (Section 7.5.2), paleoclimate evidence (Section 7.5.3), emergent constraints (Section 7.5.4), and a synthesis of all lines of evidence (Section 7.5.5). In AR5, these lines of evidence were not explicitly combined in the assessment of climate sensitivity, but as demonstrated by Sherwood et al. (2020) their combination

Sherwood et al (2020) ‘An Assessment of Earth's Climate Sensitivity Using Multiple Lines of Evidence’,⁸⁹ which influenced the IPCC’s recent estimates, though the IPCC’s own process was more qualitative.⁹⁰

The heavy tailed distributions found in some of Weitzman’s earlier work are a product of the fact that they update from a uniform prior across climate sensitivity,⁹¹ and do not systematically update on multiple lines of evidence.

I discuss this paper in [Appendix 1](#).

2.2.3. Transient Climate Response to Cumulative Emissions

The IPCC’s uncertainty about the TCRE has also narrowed. In the 2013-14 IPCC report, the 66% confidence range for the transient climate response to cumulative emissions was 0.8°C to 2.5°C per trillion tonnes of carbon.⁹² In the latest IPCC report, this has narrowed to 1.0°C to 2.3°C per trillion tonnes of carbon.⁹³ The IPCC does not give a 5% to 95% range for the transient climate response to cumulative emissions.

narrows the uncertainty ranges of ECS compared to that assessed in AR5.” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 7.5.

⁸⁹ S. Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’, *Reviews of Geophysics*, 2020.

⁹⁰ Dan Lunt, personal correspondence 9 May 2020.

See also “The broader evidence base presented in this Report and the general agreement among different lines of evidence means that they can be combined to yield a narrower range of ECS values. This can be done formally using Bayesian statistics, though such a process is complex and involves formulating likelihoods and priors (Annan and Hargreaves, 2006; Stevens et al., 2016; Sherwood et al., 2020). However, it can be understood that if two lines of independent evidence each give a low probability of an outcome being true, for example, that ECS is less than 2.0°C, then the combined probability that ECS is less than 2.0°C is lower than that of either line of evidence. On the contrary, if one line of evidence is unable to rule out an outcome, but another is able to assign a low probability, then there is a low probability that the outcome is true (Stevens et al., 2016). This general principle applies even when there is some dependency between the lines of evidence (Sherwood et al., 2020), for instance between historical energy budget constraints (Section 7.5.2.1) and those emergent constraints that use the historically observed global warming (Section 7.5.4.1). Even in this case the combined constraint will be closer to the narrowest range associated with the individual lines of evidence.” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 7.5.5.

⁹¹ On this, see also J. D. Annan and J. C. Hargreaves, ‘On the Generation and Interpretation of Probabilistic Estimates of Climate Sensitivity’, *Climatic Change* 104, no. 3–4 (1 February 2011): 423–36, <https://doi.org/10.1007/s10584-009-9715-y>.

⁹² “The transient climate response to cumulative carbon emission (TCRE) is likely between 0.8°C to 2.5°C per 1000 PgC (high confidence), for cumulative carbon emissions less than about 2000 PgC until the time at which temperatures peak” IPCC, *Climate Change 2013: The Physical Science Basis*, Fifth Assessment Report (Cambridge University Press, 2013), Technical Summary p. 81.

⁹³ “In the literature, units of °C per 1000 PgC are used, and the AR6 reports the TCRE likely range as 1.0°C to 2.3°C per 1000 PgC in the underlying report, with a best estimate of 1.65°C.” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, note 41.

2.3. Warming on different emissions scenarios

2.3.1. Current policy

It is not all that easy to estimate how much warming we might get from a given amount of cumulative CO₂ emissions. We not only need to predict CO₂ emissions from fossil fuel and industry, but also from deforestation and other forms and land use change, and in addition we need to account for non-CO₂ greenhouse gases like methane. One thing we can do is to start with projections of which representative emissions pathways we might follow and then use the IPCC's estimate of how much warming we get on those pathways.

The table below shows the warming we will get on different emissions scenarios. In Chapter 1 I argued that RCP4.5 is now plausibly the most likely scenario on current policy, and that the most likely future emissions scenario given strengthening policy is now between RCP2.6 and RCP4.5.

Table SPM.1 | Changes in global surface temperature, which are assessed based on multiple lines of evidence, for selected 20-year time periods and the five illustrative emissions scenarios considered. Temperature differences relative to the average global surface temperature of the period 1850–1900 are reported in °C. This includes the revised assessment of observed historical warming for the AR5 reference period 1986–2005, which in AR6 is higher by 0.08 [–0.01 to +0.12] °C than in AR5 (see footnote 10). Changes relative to the recent reference period 1995–2014 may be calculated approximately by subtracting 0.85°C, the best estimate of the observed warming from 1850–1900 to 1995–2014. [Cross-Chapter Box 2.3, 4.3, 4.4, Cross-Section Box TS.1]

Scenario	Near term, 2021–2040		Mid-term, 2041–2060		Long term, 2081–2100	
	Best estimate (°C)	Very likely range (°C)	Best estimate (°C)	Very likely range (°C)	Best estimate (°C)	Very likely range (°C)
SSP1-1.9	1.5	1.2 to 1.7	1.6	1.2 to 2.0	1.4	1.0 to 1.8
SSP1-2.6	1.5	1.2 to 1.8	1.7	1.3 to 2.2	1.8	1.3 to 2.4
SSP2-4.5	1.5	1.2 to 1.8	2.0	1.6 to 2.5	2.7	2.1 to 3.5
SSP3-7.0	1.5	1.2 to 1.8	2.1	1.7 to 2.6	3.6	2.8 to 4.6
SSP5-8.5	1.6	1.3 to 1.9	2.4	1.9 to 3.0	4.4	3.3 to 5.7

Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, Fig. SPM. 1.

(The 'very likely range' is the 5% to 95% range).

On current policy (i.e. RCP4.5), by 2100 we are most likely to end up with around 2.7°C of warming and there is a 1 in 20 chance of more than 3.5°C. The IPCC does not say what the chance of more than 6°C of warming is on RCP4.5, but it seems likely well below 1%, given that the upper 95th percentile is 3.5°C. This is an order of magnitude lower than estimated by Wagner and Weitzman in 2015. Since climate policy seems likely to strengthen in the future, the most likely level of warming will plausibly be 2°C to 2.5°C and the upper 95th percentile bound is probably around 3°C.

Thus, the tail risks of climate change are now much lower than they once were.

2.3.2. All fossil fuels without underground coal gasification

I argued in section 1.5 that it is difficult to come up with plausible scenarios on which we burn all of the recoverable fossil fuels. A more plausible worst-case is one in which we burn

through all of the fossil fuels that we can access without underground coal gasification. On this scenario, we would emit a further 1,200 GtC.

My calculations, using the IPCC's estimate of the TCRE imply a median warming of 3.5°C and a 66th percentile of 4.5°C. This is broadly consistent with the IPCC's own estimates. For comparison, on RCP7, total cumulative emissions are around 1,500 GtC. According to the IPCC, this implies a most likely level of warming of 3.6°C and a 95th percentile of 4.6°C.

2.3.3. The worst-case scenario

If we burned all the fossil fuels, there would most likely be around 7°C of warming relative to the pre-industrial period, and a 1 in 6 chance of warming of 9.6°C. Calculations are in this [sheet](#).

It is important to consider how long it would take us to burn through all of the fossil fuels and produce such extreme levels of warming. If there is an AI-driven explosion in carbon-intensive economic growth, then we could burn through all of the fossil fuels in the 21st Century. If there is indefinitely stalled decarbonisation, it would take several centuries to reach these extreme levels of warming.

2.4. Structural uncertainty

Until recently, it would have been reasonable to claim that there was deep structural uncertainty about equilibrium climate sensitivity. In the IPCC's 2013-14 Fifth Assessment Report, the IPCC's estimate of climate sensitivity was determined in an informal way using evidence from a range of models and from paleoclimate data.⁹⁴ The range of models used were not trying to sample the tails of the distribution. Rather, each model was trying to make a best estimate and then the distribution across these best estimates was used to estimate the final probability distribution across climate sensitivity.⁹⁵ This process is likely to underestimate tail risk.

⁹⁴ "As a consequence, in this chapter, statements using the calibrated uncertainty language are a result of the expert judgement of the authors, combining assessed literature results with an evaluation of models demonstrated ability (or lack thereof) in simulating the relevant processes (see Chapter 9) and model consensus (or lack thereof) over future projections" IPCC, *Climate Change 2013: The Physical Science Basis*, Fifth Assessment Report (Cambridge University Press, 2013), p. 1040.

⁹⁵ "Ensembles like CMIP5 do not represent a systematically sampled family of models but rely on self-selection by the modelling groups. This opportunistic nature of MMEs has been discussed, for example, in Tebaldi and Knutti (2007) and Knutti et al. (2010a). These ensembles are therefore not designed to explore uncertainty in a coordinated manner, and the range of their results cannot be straightforwardly interpreted as an exhaustive range of plausible outcomes, even if some studies have shown how they appear to behave as well calibrated probabilistic forecasts for some large-scale quantities (Annan and Hargreaves, 2010). Other studies have argued instead that the tail of distributions is by construction undersampled (Räisänen, 2007). In general, the difficulty in producing quantitative estimates of uncertainty based on multiple model output originates in their peculiarities as a statistical sample, neither random nor systematic, with possible dependencies among the members (Jun et al., 2008; Masson and Knutti, 2011; Pennell and Reichler, 2011; Knutti et al., 2013) and of spurious nature, that is, often counting among their members models with different degrees of complexities (different number of processes explicitly represented or parameterized) even within the category of general circulation models." IPCC, *Climate Change 2013: The Physical Science Basis*, p. 1040.

However, for the IPCC’s Sixth Assessment Report, concerns about model uncertainty have been greatly reduced. This is because the new estimate of climate sensitivity, builds on the work of Sherwood et al (2020) and incorporates multiple different lines of evidence, which are in part independent, including emergent constraints, the historical observational record, process-based estimates, and paleoclimate data.⁹⁶ One of the great virtues of Sherwood et al (2020) is that it carries out a sensitivity analysis of their conclusions, which is discussed in section 7 of their paper. The chart below shows the effect of:

- Different priors
- Different distributions across different lines of evidence.
- Excluding entire lines of evidence, such as historical evidence and evidence from paleoclimatic periods.

The implications of this sensitivity analysis are shown below:

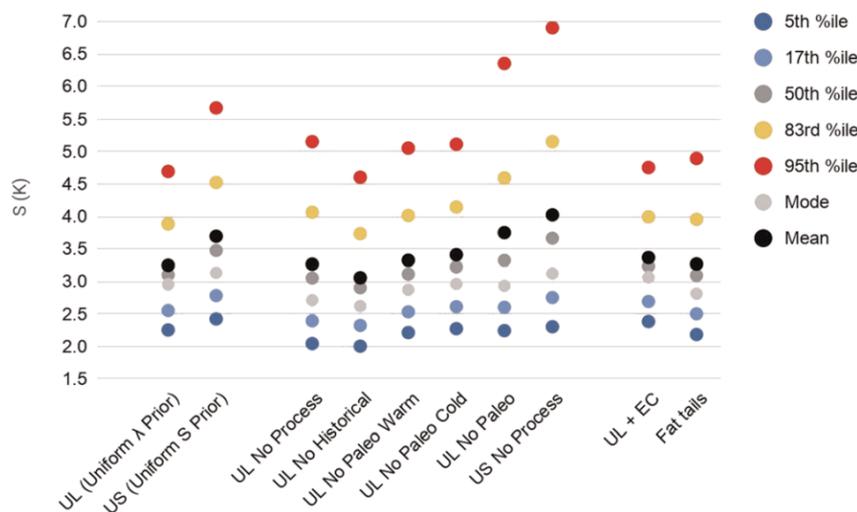


Figure 22. Graphical summary of statistics of posterior PDFs for S . UL is the Baseline calculation with a uniform prior on λ and US has a uniform prior on S . The middle group shows the effect of removing various lines of evidence in turn. UL + EC shows the impact of including the effect of emergent constraints. The effect of substituting fat-tailed distributions for some lines of evidence is also shown for the Baseline case.

Source: Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’.

Sherwood et al (2020) argue that these sensitivity tests either make little difference to their conclusions, or are implausible.⁹⁷ For example, if we exclude modern observations and models of feedback processes and assume a uniform prior across effective climate sensitivity, then the 5% to 95% range extends from 2.3°C to 6.9°C. But these assumptions are difficult to justify. For instance, for excluding the process evidence to be valid, “new

⁹⁶ S. Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’, *Reviews of Geophysics*, 2020, e2019RG000678.

⁹⁷ Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’, sec. 7.

evidence would need to come to light that would justify complete dismissal of all of the multiple elements of the process evidence (and hence much of our physical understanding of the climate system).⁹⁸

As the IPCC says in the *Sixth Assessment Report*:

“In the climate sciences, there are often good reasons to consider representing deep uncertainty, or what is sometimes referred to as unknown unknowns. This is natural in a field that considers a system that is both complex and at the same time challenging to observe. For instance, since emergent constraints represent a relatively new line of evidence, important feedback mechanisms may be biased in process-level understanding, pattern effects and aerosol cooling may be large and paleo evidence inherently builds on indirect and incomplete evidence of past climate states, there certainly can be valid reasons to add uncertainty to the ranges assessed on individual lines of evidence. This has indeed been addressed throughout Sections 7.5.1–7.5.4. Since it is neither probable that all lines of evidence assessed here are collectively biased nor is the assessment sensitive to single lines of evidence, deep uncertainty it is not considered as necessary to frame the combined assessment of ECS⁹⁹

In short, it is unlikely that all of the lines of evidence are systematically biased in one direction. As a result, according to the IPCC, structural uncertainty is now a small fraction of total uncertainty.

However, there are a number of caveats to this. Firstly, Sherwood et al (2020) only considers the possible effects of two doublings of CO₂, which would take us to around 1,100ppm.¹⁰⁰ But, on the worst-case emissions scenario, concentrations could reach 1,600ppm. The further that we push concentrations out of sample, the greater is our model uncertainty. Indeed, the IPCC says that it is confident that temperatures are proportionate to cumulative emissions up to 1 trillion tonnes of carbon; but only has medium confidence that the relationship holds for 1.5 to 3 trillion tonnes of carbon.¹⁰¹

Secondly, although the IPCC represents the broad scientific consensus on climate change, the scientific community is not unanimous in accepting that uncertainty about climate

⁹⁸ Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’, p. 69.

⁹⁹ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 7.5.5.

¹⁰⁰ “In choosing the reference scenario to define sensitivity for this assessment, for practical reasons we depart from the traditional Charney ECS definition (equilibrium response with ice sheets and vegetation assumed fixed) in favor of a comparable and widely used, so-called “effective climate sensitivity” S derived from system behavior during the first 150 years following a (hypothetical) sudden quadrupling of CO₂.” Sherwood et al., ‘An Assessment of Earth’s Climate Sensitivity Using Multiple Lines of Evidence’.

¹⁰¹ “Overall, there is high agreement between multiple lines of evidence (robust evidence) resulting in high confidence that TCRE remains constant for the domain of increasing cumulative CO₂ emissions until at least 1500 PgC, with medium confidence of it remaining constant up to 3000 PgC because of less agreement across available lines of evidence.” IPCC, *Climate Change 2021: The Physical Science Basis*, Ch. 5, sec. 5.5.1.2.1.

sensitivity has declined. Some contend that structural uncertainty remains.¹⁰² Indeed, the history of science suggests that we should probably expect scientists to be overconfident rather than underconfident, especially about complex systems.

Thirdly, the foregoing argument only applies to the IPCC's estimate of warming, it is not true for all aspects of the climate system that we might care about. For example, as discussed later in this report, there is substantial model uncertainty about future sea level rise and changing precipitation patterns, each of which could have important humanitarian consequences.

It is clear that our understanding of climate science is still imperfect, although hugely advanced over the last 50 years. Some argue that this undermines the case for climate action.¹⁰³ It is worth pausing to understand where this goes wrong. While we may be most likely to end up with 2.5°C of warming, because we are uncertain, warming could well be much higher or lower than this. But the import of this is asymmetric - lower warming might leave us with an OK outcome that is about the same as the world today, but high warming could leave us with a very bad outcome. Uncertainty makes the case for action even stronger, not weaker.

Suppose you were told that the science of rocket safety is unsettled. There is a decent chance that your rocket will burst into flames when you turn it on, but also a decent chance that it will be fine. Would this make you more or less worried about getting into a rocket?

¹⁰² See for example Joel Katzav et al., 'On the Appropriate and Inappropriate Uses of Probability Distributions in Climate Projections and Some Alternatives', *Climatic Change* 169, no. 1 (25 November 2021): 15, <https://doi.org/10.1007/s10584-021-03267-x>.

¹⁰³ Steven E. Koonin, 'Climate Science Is Not Settled', *Wall Street Journal*, 19 September 2014, sec. Life and Style, <https://online.wsj.com/articles/climate-science-is-not-settled-1411143565>.

3. Our future in paleoclimatic context

The Earth's climate has changed dramatically since the evolution of complex life at the Cambrian Explosion 540 million years ago. Putting future anthropogenic changes in context helps us to understand how bad the impact of climate change might be.

3.1. Our climate

Human influence on the climate precedes the Industrial Revolution. In the thousands of years prior to 1800 AD, much of the world was deforested for agriculture, releasing greenhouse gases into the atmosphere. Pre-industrial deforestation, along with methane-producing rice farming, released greenhouse gases into the atmosphere,¹⁰⁴ warming the planet today by a few tenths of a degree Celsius.¹⁰⁵ There is disagreement in the literature about the size of pre-industrial emissions.

Nonetheless, most of our species' influence on the climate has come from the burning of fossil fuels since the Industrial Revolution. Immediately prior to the Industrial Revolution, CO₂ concentrations were 278ppm, and global average surface temperature was around 14°C.¹⁰⁶ Today, CO₂ concentrations are at 410ppm and the global average temperature is around 1 degree higher.

3.2. Hothouses and rapid warming

To understand the risks to sentient life from global warming, it is useful to compare the anthropogenic future to the distant past. The Earth's climate has changed dramatically over

¹⁰⁴ This includes methane and carbon dioxide. Estimates of the size of the emissions differ. IPCC, *Climate Change: The Physical Science Basis*, sec 6.2.2; W. F. Ruddiman et al., 'The Early Anthropogenic Hypothesis: A Review', *Quaternary Science Reviews* 240 (15 July 2020): 1, <https://doi.org/10.1016/j.quascirev.2020.106386>.

¹⁰⁵ Estimates of the size of the effect vary. On the higher estimates of carbon dioxide emissions, pre-industrial land use change caused 0.76C of warming. See Feng He et al., 'Simulating Global and Local Surface Temperature Changes Due to Holocene Anthropogenic Land Cover Change', *Geophysical Research Letters* 41, no. 2 (2014): 623–31, <https://doi.org/10.1002/2013GL058085>. On the lower estimates, increasing carbon dioxide concentrations by 10ppm would have caused warming of 0.16C (assuming a climate sensitivity of 3C). The IPCC estimates that there is a 33% to 66% chance that anthropogenic land use change is responsible for the increase in methane 5,000 years ago. Ruddiman et al (2003) estimates that the methane release caused 0.25C of warming. Following the IPCC, assuming that there is a 50% chance that this was anthropogenic, the expected warming from anthropogenic methane release was 0.13C. Overall, this suggests that the combined effect of pre-industrial anthropogenic carbon dioxide and methane was warming of around 0.3C. IPCC, *Fifth Assessment Report, Climate Change: The Physical Science Basis*, sec 6.2.2.2; William F. Ruddiman, 'The Anthropogenic Greenhouse Era Began Thousands of Years Ago', *Climatic Change* 61, no. 3 (2003): 285.

¹⁰⁶ "Quantitative comparison with fully coupled climate model simulations indicates that global average temperatures were about 29, 26, 23 and 19 degrees Celsius in the early, early middle, late middle and late Eocene, respectively, compared to the preindustrial temperature of 14.4 degrees Celsius." Margot J. Cramwinckel et al., 'Synchronous Tropical and Polar Temperature Evolution in the Eocene', *Nature* 559, no. 7714 (July 2018): 382–86, <https://doi.org/10.1038/s41586-018-0272-2>.

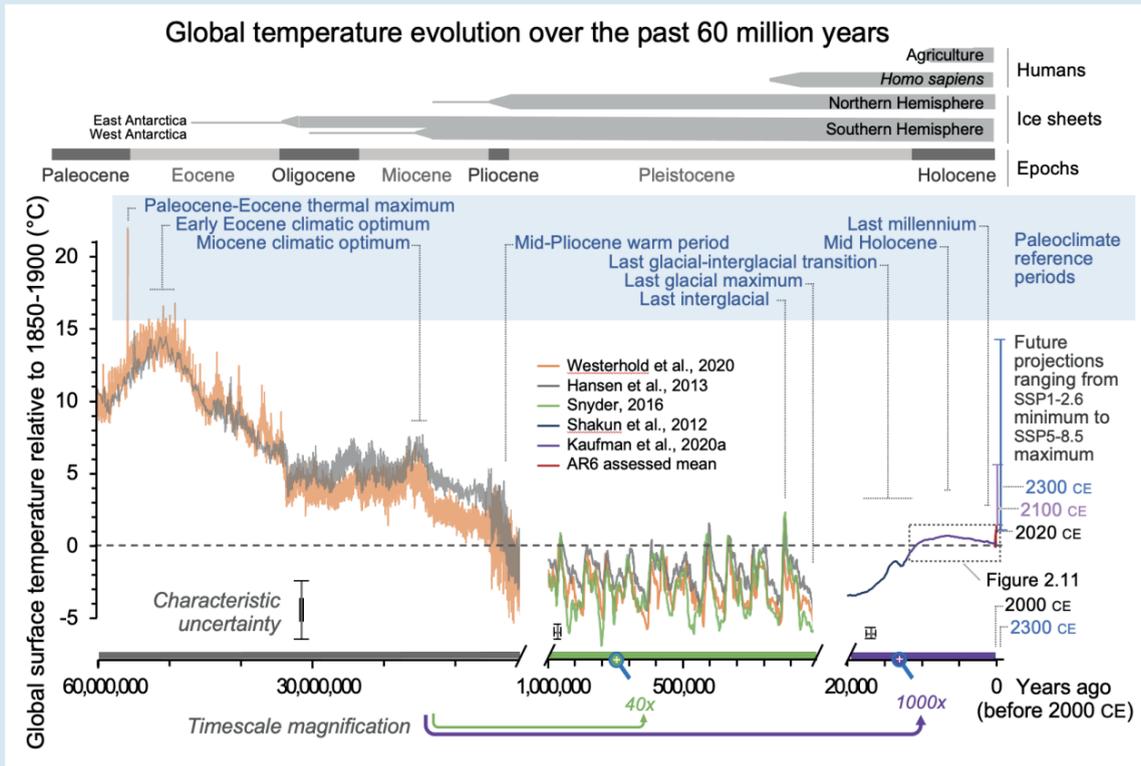
the last half a billion years, shifting from extremely cold to extremely hot. Better understanding our climatic past is crucial for understanding two important issues:

1. **Ecosystem trauma.** How did ecosystems cope with much hotter climates and with rapid warming? If past warm periods caused substantial loss of animal, plant and insect life, then that might indicate that global warming will make life harder for sentient life in the future.
2. **Tipping points.** What is the risk that we will set off tipping points if we pass certain climatic thresholds?

I will discuss ecological trauma in this Chapter and tipping points in Chapter 8.

In this Chapter, my main focus is on the relevance of past climate change for *human* flourishing in the face of future climate change. I discuss the potential impact of future climate change for ecosystems in Chapter 5.

The chart below from the 2021 IPCC Sixth Assessment report shows future potential warming in the context of the last 60 million years. For reference, on current policy, global temperatures are most likely to increase by 2.5°C relative to pre-industrial times. On RCP8.5, temperature would increase by 5°C by 2100.



Cross-Chapter Box 2.1, Figure 1 | Global mean surface temperature (GMST) over the past 60 million years (60 Myr) relative to 1850–1900 shown on three time scales. Information about each of the nine paleo reference periods (blue font) and sections in AR6 that discuss these periods are listed in Cross-Chapter Box 2.1 Table 1. Grey horizontal bars at the top mark important events. Characteristic uncertainties are based on expert judgement and are representative of the approximate midpoint of their respective time scales; uncertainties decrease forward in time. GMST estimates for most paleo reference periods (Figure 2.34) overlap with this reconstruction, but take into account multiple lines of evidence. Future projections span the range of global surface air temperature best estimates for SSP1–2.6 and SSP5–8.5 scenarios described in Section 1.6. Range shown for 2100 is based on CMIP6 multi-model mean for 2081–2100 from Table 4.5; range for 2300 is based upon an emulator and taken from Table 4.9. Further details on data sources and processing are available in the chapter data table (Table 2.SM.1).

Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, Cross-Chapter Box 2.1, Fig. 1.

The chart below from Scotese et al (2021) shows global average temperatures further back in time:

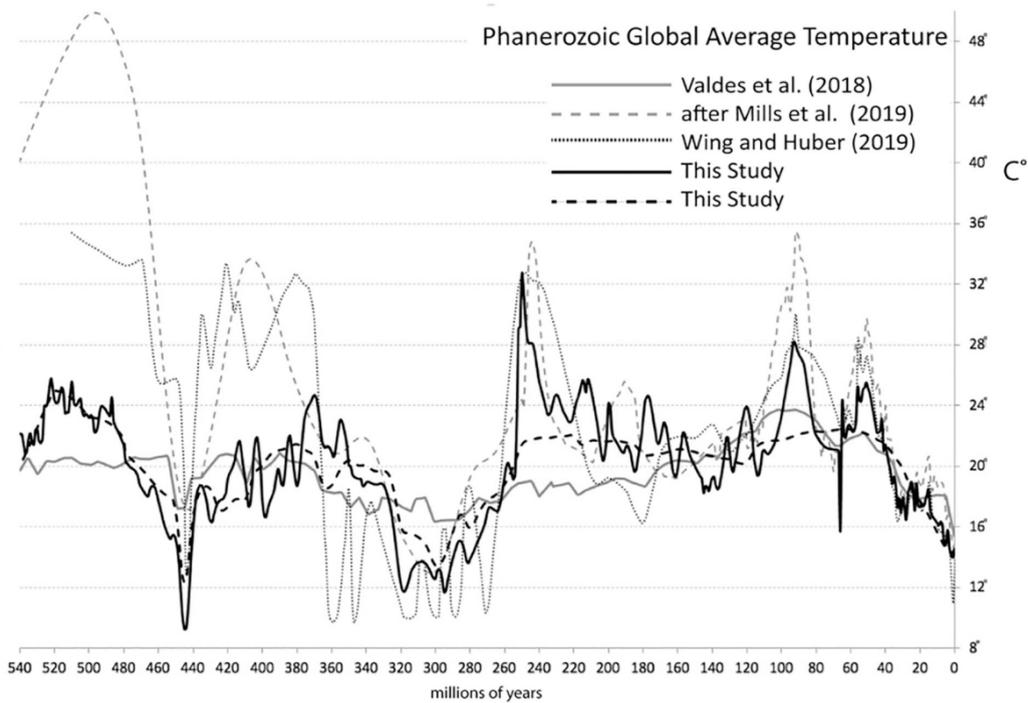
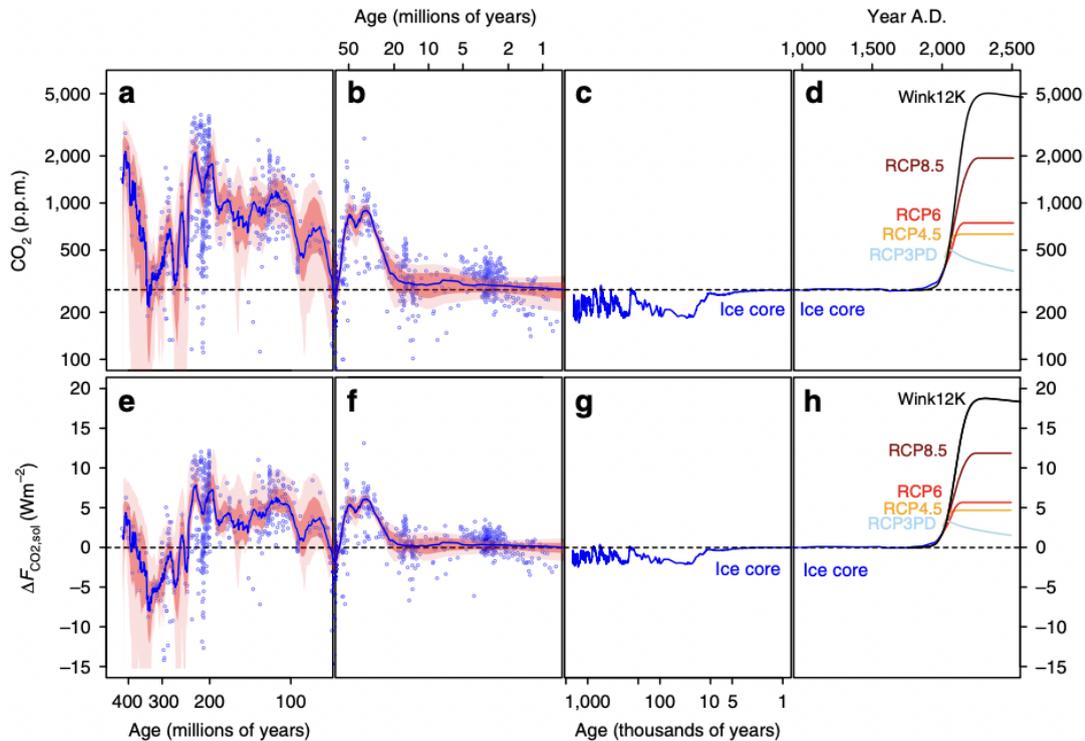


Fig. 1. Estimates of Phanerozoic Global Average Temperature (GAT). Sources: Wing and Huber (2019), Valdes et al. (2018), Mills et al. (2019), and This study.

Source: Christopher R. Scotese et al., 'Phanerozoic Paleotemperatures: The Earth's Changing Climate during the Last 540 Million Years', *Earth-Science Reviews* 215 (1 April 2021): fig. 1, <https://doi.org/10.1016/j.earscirev.2021.103503>.

The chart below from Foster et al (2017) shows the evolution of CO₂ concentrations (top pane) and the warming effect of CO₂ (bottom pane) over the last half billion years.¹⁰⁷

¹⁰⁷ Gavin L. Foster, Dana L. Royer, and Daniel J. Lunt, 'Future Climate Forcing Potentially without Precedent in the Last 420 Million Years', *Nature Communications* 8, no. 1 (4 April 2017): 14845, <https://doi.org/10.1038/ncomms14845>.



Source: Gavin L. Foster, Dana L. Royer, and Daniel J. Lunt, 'Future Climate Forcing Potentially without Precedent in the Last 420 Million Years', *Nature Communications* 8, no. 1 (4 April 2017): 14845, <https://doi.org/10.1038/ncomms14845>, Fig. 4.

The table below shows key facts about the different periods I will discuss in this Chapter.

Period	When?	Temp vs pre-industrial	Global warming per century	Regional warming per century	CO ₂ ppm
RCP8.5	2100	5°C	4°C	4-10°C	900ppm
Current policy	2100	2.5°C	1.5°C	2.5-5°C	600ppm
Transition from last ice age to Holocene	20k to 6k years ago			2-15°C	
Last interglacial	127k years ago	1°C			280ppm
Mid-Pliocene Warm Period	3m years ago	3°C			390ppm
Miocene Climatic Optimum	16 million years ago	7°C			500ppm
Eocene Climatic Optimum	50m years ago	14°C			1,800ppm
Paleocene-Eocene Thermal Maximum	56m years ago	17°C	0.05°C		2,300ppm
Mid-Cretaceous	90m years ago	20°C			1,000ppm

Permian to early Triassic	250m years ago	17°C			3,000ppm
---------------------------	----------------	------	--	--	----------

(Blanks mean that no good data is available, or that the category is not applicable)

Transitions between geological periods are often caused by or associated with climatic changes. For example, the transition from the Permian to the Triassic was associated with the eruption of the Siberian Traps and the huge global warming and extinction that followed. Transitions into and out of glacial periods are primarily determined by the Earth's orbit in relation to the Sun, and the tilt of the Earth on its axis.¹⁰⁸

To estimate past CO₂ concentrations and temperatures, we have to rely on imperfect proxies, which creates substantial uncertainty.

3.3. Were past greenhouses inhospitable?

The picture that emerges from the paleoclimate is as follows:

- Before 180 million years ago, climate change was a probable cause of massive loss of biodiversity
- From 180 million years ago until today, climate change has not been associated with elevated species loss.

In the remainder of this section, I will discuss episodes in these paleoclimatic periods in depth, and try to get to the bottom of why the response to warming was so different in these different geological periods.

3.3.1. Pre-Cretaceous warming and mass extinctions

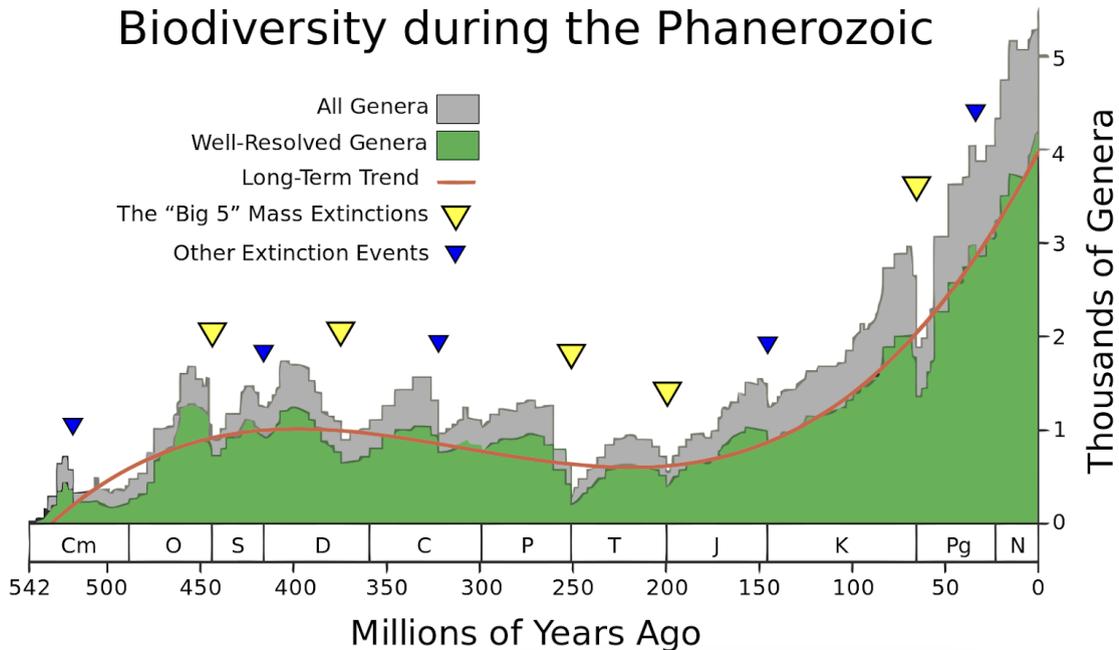
Before I discuss trends in hospitability, we should first get clear on some taxonomic terms that I will use in what follows. *Genus* is a taxonomic rank that comes above species. For example, *homo sapiens* are a species in the *homo* genus, which once included other *homo* species, such as neanderthals (*homo neanderthalensis*). Lions and jaguars are two species in the genus *panthera*.

This [chart from Wikipedia](#) using data from Rohde and Muller (2005) shows the trend in biodiversity during our current eon, the Phanerozoic, from the Cambrian explosion, when almost all animal phyla appeared, until today.¹⁰⁹

¹⁰⁸ IPCC, Climate Change: The Physical Science Basis, 399.

¹⁰⁹ Robert A. Rohde and Richard A. Muller, 'Cycles in Fossil Diversity', *Nature* 434, no. 7030 (March 2005): 208–10, <https://doi.org/10.1038/nature03339>. For a similar picture for marine biodiversity see Jeremy B. C. Jackson and Kenneth G. Johnson, 'Measuring Past Biodiversity', *Science* 293, no. 5539 (28 September 2001): 2401–4, <https://doi.org/10.1126/science.1063789>.

Biodiversity during the Phanerozoic



145 million years ago, the Jurassic (J on the chart above) ended and the Cretaceous (K) began. Prior to the Cretaceous, biodiversity was low and major extinction events occurred with striking regularity. Since then, with the exception of the extinction event that killed off the dinosaurs, it has been relatively plain sailing for Earth's various species, until humans started killing off other species themselves.

Mass extinctions are categorised as times when the Earth loses >75% of its species in the space of 2 million years or less.¹¹⁰ In the last 540 million years, there have been five mass extinctions.

1. Ordovician–Silurian: 450–440 million years ago.
2. Late Devonian: 375–360 million years ago.
3. Permian–Triassic (End Permian): 252 million years ago.
4. End Triassic: 201.3 million years ago.
5. End Cretaceous: 66 million years ago.

Aside from these events, there have been many other major extinction events that do not qualify as mass extinctions.

Scholars disagree about the causes of mass extinctions, but the most popular explanation of the causes of these and other extinction events is volcanic eruptions, apart from the extinction of the dinosaurs, which is widely agreed to have been due to an asteroid impact.

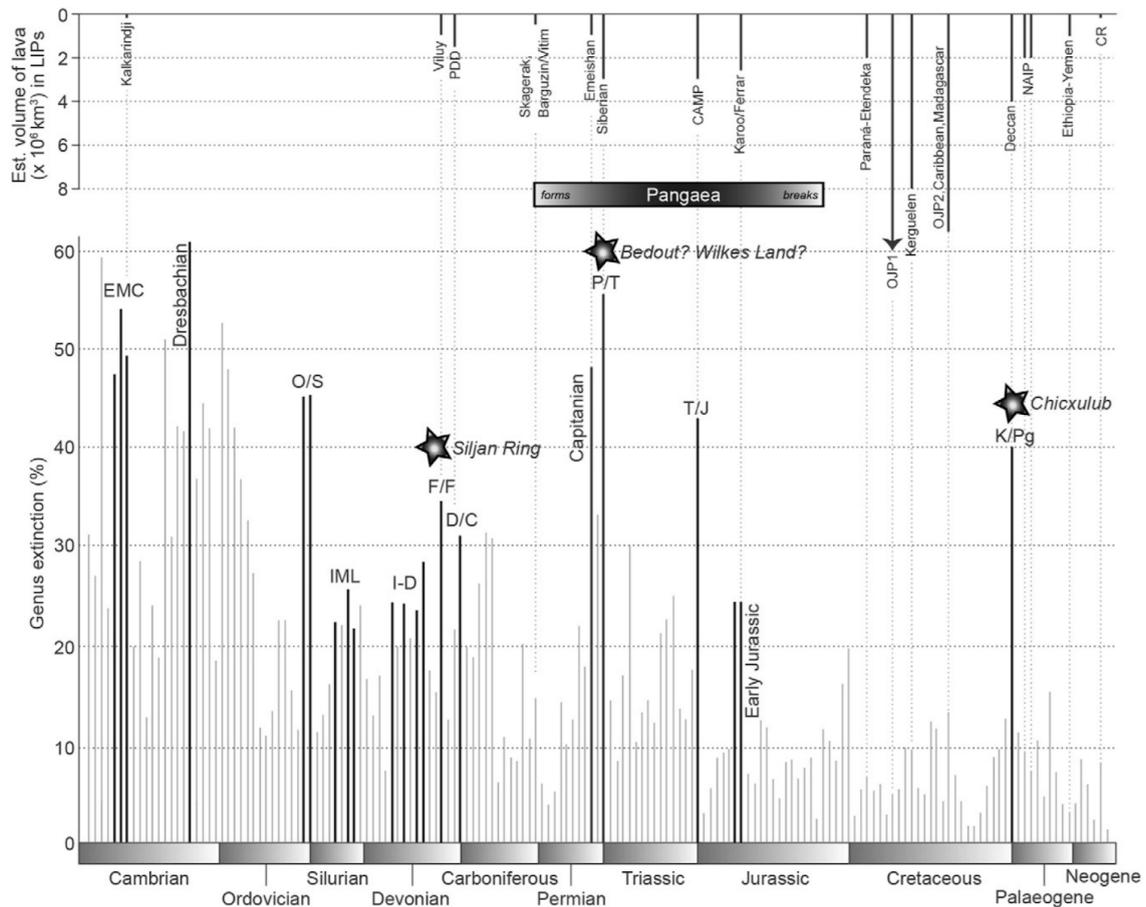
¹¹⁰ Anthony D. Barnosky et al., 'Has the Earth's Sixth Mass Extinction Already Arrived?', *Nature* 471 (2 March 2011): 51.

Table 1
Summary of data and proposed causal mechanisms implicated in mass extinctions since the Early Cambrian.

Extinction (age)	Associated LIP	Associated impact structure	Global warming or cooling?	Ocean acidification?	Marine anoxia?	Carbon isotope shift	Notes and other postulated causes
Early/Middle Cambrian (Botomian)	Kalkarindji	None	?	?	Yes	-4‰	ROECE Event in carbon isotope stratigraphy. Maybe several extinction events – poorly constrained.
Dresbachian	None	None	Warming?	?	Yes	+5‰	SPICE Event in carbon isotope stratigraphy. Gamma-ray burst?
End Ordovician	Speculated	None	Cooling (phase 1) and warming (phase 2)	?	Partly	+7‰ followed by -7‰	
Ireviken Event	None	None	Cooling?	?	Yes	-4‰ imposed on a positive trend	
Mulde Event	None	None	Cooling?	?	Yes	+4‰	Starvation amongst planktonic larvae driven by severe drop in primary planktonic productivity
Lau Event (Ludfordian)	None	None	Cooling?	?	Yes	+6‰	
Kačák Event (Eifellian)	None	None	Warming?	?	Yes	+2‰	
Thaganic Event (Givetian)	None	None	Warming?	?	Yes	+2‰	
Frasnian-Famennian	Viluy Traps, PDD?	Siljan Ring?	Warming (+9 °C) imposed on cooling pulses	?	Yes	up to +4‰	
Hangenberg Event (End Devonian)	PDD?	Woodleigh, Western Australia?	Warming and cooling (including glaciation)	?	Yes	up to +6‰	A prolonged and diachronous extinction (several 100kyr) on par with the "big 5" but poorly understood
Capitanian	Emeishan Traps	None	? both have been invoked	Possibly	Yes (only regionally)	-6‰ (in China)	Volcanic darkness and photosynthetic shutdown; toxic metal (Hg) poisoning
End Permian	Siberian Traps	Bedout? Wilkes Land?	Warming (+10 °C)	Probably	Yes	up to -8‰	Acid rain; toxic metal (Hg) poisoning; UV-B damage
Smithian/Spathian	Siberian Traps (late stages)	None	Warming (+6 °C)	?	Yes	-6‰ followed by +6‰	Toxic metal (Hg) poisoning
Carnian	Wrangellia	None	Warming (+7 °C)	?	Yes	-5‰	Major radiations as well as extinctions
End Triassic	CAMP	None	Warming (+6 °C)	Probably	?	-5‰	Seismite at the extinction level: CAMP or bolide? Toxic metal (Hg) poisoning
Early Jurassic	Karoo/Ferrar	None	Warming (+7 °C)	?	Yes	-7‰ in $\delta^{13}C_{org}$ -3‰ in $\delta^{13}C_{carb}$	Toxic metal (Hg) poisoning
End Cretaceous	Deccan Traps	Chicxulub	Warming (+4 °C)	?	No	-2‰	Toxic metal (Hg) poisoning

Source: David P. G. Bond and Stephen E. Grasby, 'On the Causes of Mass Extinctions', *Palaeogeography, Palaeoclimatology, Palaeoecology*, 478 (15 July 2017): 3–29, <https://doi.org/10.1016/j.palaeo.2016.11.005>.

This figure shows the correlation between volcanism and major extinction events. It also shows how extinctions were especially bad during the time of the Pangean supercontinent, a point I return to at the end of this chapter.



Source: David P. G. Bond and Stephen E. Grasby, 'On the Causes of Mass Extinctions', *Palaeogeography, Palaeoclimatology, Palaeoecology*, 478 (15 July 2017): 3–29, <https://doi.org/10.1016/j.palaeo.2016.11.005>.

The lines down from the top show the volume of lava from different volcanic eruptions, while the lines from the bottom show the percentage of genera killed. In '[On the causes of mass extinctions](#)', Bond and Grasby note

“The temporal link between large igneous province (LIP) eruptions and at least half of the major extinctions of the Phanerozoic implies that large scale volcanism is the main driver of mass extinction.”¹¹¹

Other posited causes for mass extinctions include Near Earth Objects and gamma ray bursts.¹¹² A huge asteroid impact in present day Mexico is currently the most popular explanation for the extinction of the dinosaurs 65 million years ago at the end of the Cretaceous, though the eruption of the Deccan Traps also coincided with that event, as shown in the chart above.¹¹³

¹¹¹ Bond and Grasby, 'On the Causes of Mass Extinctions', p. 3. Large Igneous Provinces (LIPs) are the igneous rock formations left over after large volcanic eruptions. The Phanerozoic is the current geological eon, which started 540 million years ago.

¹¹² Bond and Grasby, 'On the Causes of Mass Extinctions'.

¹¹³ Alfio Alessandro Chiarenza et al., 'Asteroid Impact, Not Volcanism, Caused the End-Cretaceous Dinosaur Extinction', *Proceedings of the National Academy of Sciences* 117, no. 29 (21 July 2020): 17084–93, <https://doi.org/10.1073/pnas.2006087117>.

The chart above also nicely illustrates trends in ecological stress. In the 80 million years from the first Permian extinction event, the Capitanian, to the early Jurassic extinction events, the average rate of global genus extinctions in extinction events is 15-20%, and 12 events produced global genus extinction rates in excess of 15%. But in the 145 million years since the end of the Jurassic, the average rate of global genus extinctions from extinction events has been around 5% and never passed 15%, except for the death of the dinosaurs.

There is no better illustration of the damage that volcanoes caused than the greatest ecological disaster of all time: the Permian-Triassic extinction.

Warming at the Permian-Triassic boundary

252 million years ago, the eruption of the Siberian Traps put 30 to 40 trillion tonnes of carbon into the atmosphere,¹¹⁴ and CO₂ concentrations rose to around 3,000ppm.¹¹⁵ Average ocean temperatures in South China before the end Permian disaster were 20°C and peaked at 40°C in the early Triassic.¹¹⁶ At no point did ocean temperatures in Southern China drop below 32°C in the 5 million years after the end Permian event. According to Wignall, these temperatures are the highest ever recorded.¹¹⁷ For comparison, modern equatorial sea surface temperatures typically average around 28°C and never exceed 30C.¹¹⁸ Peak ocean temperatures were the same as you would find in a bowl of very hot soup.

Bond and Grasby (2017) comment

“With 90% marine species loss, widespread devastation on land including the only recorded mass extinction of insects it is Earth's greatest ever biotic crisis.”¹¹⁹

Many volcanic kill mechanisms have been proposed for the Permian-Triassic extinction. Many of these mechanisms would also have been in play in the other Pangean extinction events.

Volcanic kill mechanisms

Volcanoes inject a wide array of gases into the atmosphere, with water vapour, CO₂ and sulphur dioxide the most important volumetrically. Volcanoes also release halogens, which

¹¹⁴ “Clarkson et al.[12] estimated that the two phases of P-Tr extinction were driven first by a small addition of isotopically light C (e.g. methane) and then by a massive addition of isotopically heavy C (e.g. from decarbonation of limestones intruded by Siberian Traps dikes and sills), with a total emission of from 30 000–40 000 Pg C (consistent with an independent estimate by Svensen et al. [9]).” Lee R. Kump, ‘Prolonged Late Permian–Early Triassic Hyperthermal: Failure of Climate Regulation?’, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, no. 2130 (13 October 2018): 20170078, <https://doi.org/10.1098/rsta.2017.0078>.

¹¹⁵ Ying Cui and Lee R. Kump, ‘Global Warming and the End-Permian Extinction Event: Proxy and Modeling Perspectives’, *Earth-Science Reviews*, Global review of the Permian-Triassic mass extinction and subsequent recovery: Part II, 149 (1 October 2015): 5–22, <https://doi.org/10.1016/j.earscirev.2014.04.007>.

¹¹⁶ P. B. Wignall, *The Worst of Times: How Life on Earth Survived Eighty Million Years of Extinctions* (Princeton: Princeton University Press, 2015), p. 97-98.

¹¹⁷ Wignall, *The Worst of Times*, p. 97-98.

¹¹⁸ Wignall, *The Worst of Times*, p. 98.

¹¹⁹ Bond and Grasby, ‘On the Causes of Mass Extinctions’, 10.

could be damaging insofar as they destroy the ozone layer, which protects plant and animal life from UV radiation.

The kill mechanisms for each of the gases are as follows:

- **Sulphur dioxide**
 - Causes volcanic darkness, cooling and photosynthetic shutdown
 - Causes acid rain, which is bad for plants and contributes to ocean anoxia
- **Carbon dioxide**
 - Leads to warming which exposes organisms to thermal stress, and changes ecosystems which creates adaptation challenges for some species.
 - Warming reduces the capacity of the ocean to absorb oxygen, which can lead to ocean anoxia, which is implicated in several marine extinction events.
 - CO₂ dissolves in the ocean causing ocean acidification, which is generally bad for shellfish and marine organisms with carbonate shells.
 - CO₂ can build up in tissues, a process known as hypercapnia, with potentially fatal consequences.
- **Gases like nitrous oxide, chlorine and fluorine**
 - Damage the ozone layer, which protects life on Earth from harmful UV rays.
 - Causes acid rain.
- **Release of toxic metals, especially mercury**
 - These toxic metals can be dangerous to life, though their role in mass extinctions has yet to be fully evaluated.

A single pulse of sulphur dioxide into the atmosphere can cause cooling for two to three years, after which it is rained out. Thus, sulphur dioxide is probably only capable of driving death-by-cooling if eruptions were frequent and of high volume and were sustained for several centuries at a time. Unfortunately, the geological record of Large Igneous Provinces is not sufficiently resolved to permit an evaluation of whether this has actually happened during a mass extinction interval.¹²⁰

CO₂ is especially damaging to marine life and can cause warming on geological timescales: about a third of CO₂ emissions remain in the atmosphere 10,000 years after it is injected.¹²¹ As we have seen, the volume of CO₂ released from some volcanic eruptions in the past is large relative to potential anthropogenic emissions.

Huge releases of CO₂ can do great damage to marine ecosystems. But extinctions on land were plausibly driven by other gases. Recent research suggests that the eruption of the Siberian Traps was so damaging because the eruption happened to punch through and evaporate salt deposits, releasing halogens causing ozone destruction, which was the major driver of extinctions of land-based communities. That is plausible why end-Permian level extinction events do not happen more often.¹²²

¹²⁰ Bond and Grasby, 'On the Causes of Mass Extinctions', 16.

¹²¹ N. S. Lord et al., 'An Impulse Response Function for the "Long Tail" of Excess Atmospheric CO₂ in an Earth System Model', *Global Biogeochemical Cycles* 30, no. 1 (2016): Fig. 2, <https://doi.org/10.1002/2014GB005074>.

¹²² Thanks to Matthew Huber for raising this point. "Volatile emissions to the atmosphere associated with the Siberian Traps eruptions at the Permian-Triassic boundary were sourced from the outgassing

CO₂ is the main volcanic kill mechanism relevant to anthropogenic climate change.¹²³ While large releases of CO₂ place great strain on marine ecosystems, they were probably not sufficient to cause land-based extinctions. Thus, although past volcanic eruptions did cause land-based species extinctions, the mechanism was probably not CO₂. There might be some other reason that future anthropogenic warming will cause land-based species extinctions, but the paleoclimate record should not update us towards that view.

Still, the evidence from this period suggests that the release of carbon and global warming could do enormous damage.

of primary magmas and the sedimentary host rocks into which they were intruded. Halogens in volcanic gases may have played an important role in environmental degradation and in stratospheric ozone destruction. Here we investigate how halogens behave during the interaction between salts and basalt magma emplaced as sills and erupted as lava. We present whole-rock, trace, and halogen concentrations for a suite of samples from three locations in the Siberian Traps Large Igneous Province, including basalt lavas erupted, and dolerites intruded into both organic-bearing shales and evaporites. Dolerites are enriched in Cl, Br, and I; their enrichment in Cl is similar to MORB and OIB that have been inferred to have assimilated seawater. The dolerites exhibit halogen compositional systematics, which extend towards both evaporites and crustal brines. Furthermore, all analyzed samples show enrichment in Rb/Nb; with the dolerites also showing enrichment in Cl/K similar to MORB and OIB that have been inferred to have assimilated seawater. We infer that samples from all three locations have assimilated fluids derived from evaporites, which are components of crustal sedimentary rocks. We show that up to 89% of the chlorine in the dolerites may have been assimilated as a consequence of the contact metamorphism of evaporites. We show, by thermal modeling, that halogen transfer may occur via assimilation of a brine phase derived from heating evaporites. Halogen assimilation from subcropping evaporites may be pervasive in the Siberian Traps Large Igneous Province and is expected to have enhanced emissions of Cl and Br into the atmosphere from both intrusive and extrusive magmatism.” Svetlana Sibik et al., ‘Halogen Enrichment of Siberian Traps Magmas During Interaction With Evaporites’, *Frontiers in Earth Science* 9 (2021), <https://www.frontiersin.org/articles/10.3389/feart.2021.741447>.

But see also “We still do not have all the answers to the questions posed by Pangean mass extinctions. The devastation of land communities is especially hard to explain. The extinction of Late Permian terrestrial communities is a truly awesome phenomenon, which might be related to atmospheric changes such as ozone destruction. Massive volcanic halogen emissions provide one cause, but then why did this only happen during Pangea’s lifetime? More recent eruptions, such as those of the NAIP, would be expected to also emit huge amounts of halogens. Alternatively, some geologists have attempted to link changes in the oceans to events on land. Thus anoxic oceans may have leaked hydrogen sulfide into the atmosphere, where it would interfere with the formation of ozone.

Unfortunately, this probably is not the answer. Hydrogen sulfide is immensely reactive with oxygen and is unlikely to ever reach the stratosphere and damage the ozone shield; instead, it will oxidize rapidly close to the sea surface. Perhaps terrestrial warming was the stress factor on land, but it would have been most harsh on forests adapted to living in cold conditions. The extinctions show no such temperature dependence—equatorial forests suffered as badly as those at high latitudes. And so the puzzle remains.” Wignall, *The Worst of Times*, 175.

¹²³ [Sulphur dioxide emissions](#) peaked in 1980 and have declined every decade since then up to 2010 (there are no data beyond then). [Mercury emissions](#) are mainly a product of low economic development, and are now low in high-income countries. Emissions of [ozone depleting substances](#) peaked in 1988 and have declined since then. Since around 2005, the ozone hole has slowly started to shrink.

3.3.2. Biodiversity after Pangea

The connection between warming and ecological stress disappeared after the break-up of the Pangea supercontinent 180 million years ago. In this section, I will discuss some periods involving high temperatures and/or rapid warming. Throughout, I will note disanalogies between these periods and future anthropogenic warming

An excellent summary of some of these periods is provided in [Willis and MacDonald \(2011\) 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World'](#)

Cretaceous warming

Period	When?	Temp vs pre-industrial	Global warming per century	Regional warming per century	CO ₂ ppm
Mid-Cretaceous (Tierney)	90m years ago	20°C			1,000ppm

In the mid-Cretaceous, around 90 million years ago, temperatures were around 20°C warmer than pre-industrial levels, while CO₂ levels were between 500ppm and 1,000ppm.¹²⁴ (The mid-Cretaceous may have been so warm in part for non-greenhouse gas-related reasons.)¹²⁵ Despite these extreme temperatures, as the charts above show, the rate of species extinction did not seem to increase. Unlike the Paleocene-Eocene Thermal Maximum, discussed below, there was no absence of plankton in tropical regions.¹²⁶

These episodes may not be relevant to our current geological era - the Cenozoic - which started after the extinction of the dinosaurs. Extinction events prior to the Cenozoic killed much of life on Earth, and the ecosystems that survived are very different to those that

¹²⁴ Jessica E. Tierney et al., 'Past Climates Inform Our Future', *Science* 370, no. 6517 (6 November 2020): fig. 1, <https://doi.org/10.1126/science.aay3701>.

¹²⁵ "Atmospheric CO₂ concentrations generally follow these swings in global temperature (Fig. 1). Geochemical modeling demonstrates that the balance of geological sources (degassing through volcanism) and sinks (weathering and sedimentation) explains the general features of CO₂'s trajectory(8) and establishes causality high CO₂ leads to high temperatures. The apparent exceptions to this rule, including the end-Cretaceous and early Paleocene (70 to 60 Ma) and the Miocene (23 to 5.3 Ma), are areas of active research. One explanation for the decoupling of CO₂ and temperature is that uncertainties associated with the proxies blur the relationship. Estimation of past CO₂ is challenging. Beyond the ice core record (9), CO₂ information comes from geochemical data, such as isotope ratios of boron and carbon, or paleobotanical indicators such as the density of leaf stomata. All of these proxies require assumptions about the physical, chemical, and biological state of the past that are not completely understood, sometimes leading to misinterpretations of the signal (10). Proxy methodologies and assumptions continue to be refined, and there is some indication that CO₂ at the end of the Cretaceous may have been higher than that shown in Fig. 1 (11). It is also possible that these discrepancies have another explanation, such as a greater than expected role for non-CO₂ forcings and feedbacks. If the paleoclimate record has taught us anything, it is that the more we probe, the more we learn." Jessica E. Tierney et al., 'Past Climates Inform Our Future', *Science* 370, no. 6517 (6 November 2020), <https://doi.org/10.1126/science.aay3701>.

¹²⁶ "No similar temporal absence of eukaryotic plankton has been observed in tropical regions during the extreme warmth of OAE2 (Oceanic Anoxic Event 2), across the Cenomanian-Turonian transition (~94 Ma), where comparable TEXH 86 temperature estimates were obtained (51)." Joost Frieling et al., 'Extreme Warmth and Heat-Stressed Plankton in the Tropics during the Paleocene-Eocene Thermal Maximum', *Science Advances* 3, no. 3 (1 March 2017): e1600891, <https://doi.org/10.1126/sciadv.1600891>.

preceded them.¹²⁷ It is therefore instructive to consider warming events that happened during the Cenozoic.

The Paleocene-Eocene Thermal Maximum and Eocene Climatic Optimum

Period	When?	Temp vs pre-industrial	Global warming per century	CO ₂ ppm
Eocene Climatic Optimum	50m years ago	14°C		1,800ppm
Paleocene-Eocene Thermal Maximum	56m years ago	17°C	0.05C	2,300ppm

During the Paleocene-Eocene Thermal Maximum, temperatures increased by around 5°C over the course of 3,000 to 20,000 years,¹²⁸ This occurred on top of a very warm background: at the peak of the PETM, temperatures were around 17C warmer than pre-industrial, with a 5% to 95% range 10°C–25°C. Equatorial sea surface temperatures passed 36C,¹²⁹ compared to around 28C today.¹³⁰

There is disagreement about the size of the carbon release, but most sources agree it is well in excess of 1.5 trillion tonnes of carbon in the form of CO₂ and/or methane,¹³¹ with one recent estimate putting emissions at around 10 trillion tonnes of carbon over 50,000 years.¹³²

¹²⁷ Thanks to Matthew Huber for raising this point.

¹²⁸ “During the PETM (56 Ma) CO₂ rapidly rose from about 900 ppm to about 2000 ppm (Anagnostou et al., 2020; Gutjahr et al., 2017; Schubert & Jahren, 2013; Table 41 2.1) in 3–20 kyr (Gutjahr et al., 2017; Turner, 2018; Zeebe et al., 2016).” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.2.3.1. “A major new compilation of proxy temperature data (Hollis et al., 2019) analysed using multiple statistical approaches (Inglis et al., 2020) indicates that GMST was 10°C–25°C (90% range) warmer than 1850–1900, or about 5°C warmer relative to the pre-PETM state. A related synthesis study also estimates that PETM warmed by 5°C (no uncertainty assigned; Zhu et al., 2019). A recent benthic isotope compilation (Westerhold et al., 2020) transformed to GMST based on the formulation by Hansen et al. (2013c) (Cross-Chapter Box 2.1, Figure 1), and adjusted to 49 1850-1900 by adding 0.36°C, shows an increase of GMST by about 10°C during the PETM.” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.3.1.1.1.

¹²⁹ “On the basis of planktonic foraminiferal Mg/Ca and oxygen isotope ratios and the molecular proxy TEXH 86, latest Paleocene equatorial SSTs were ~33°C, and TEXH 86 indicates that SSTs rose to >36°C during the PETM.” Joost Frieling et al., ‘Extreme Warmth and Heat-Stressed Plankton in the Tropics during the Paleocene-Eocene Thermal Maximum’, *Science Advances* 3, no. 3 (1 March 2017): e1600891, <https://doi.org/10.1126/sciadv.1600891>.

¹³⁰ Yadong Sun et al., ‘Lethally Hot Temperatures During the Early Triassic Greenhouse’, *Science* 338, no. 6105 (19 October 2012): Fig. 2, <https://doi.org/10.1126/science.1224126>.

¹³¹ “This pattern is best explained by massive (>>1500 Gt) carbon input from at least one but likely multiple reservoirs in the shape of CO₂ and/or CH₄” Joost Frieling et al., ‘Tropical Atlantic Climate and Ecosystem Regime Shifts during the Paleocene–Eocene Thermal Maximum’, *Climate of the Past* 14, no. 1 (15 January 2018): 39–55, <https://doi.org/10.5194/cp-14-39-2018>.

¹³² “Using our preferred age model (R07sm; Extended Data Table 1a), we diagnose a cumulative PETM carbon release that reaches about 10,200 Pg, with almost all emissions occurring in the first 50 kyr (Fig. 3d). Marcus Gutjahr et al., ‘Very Large Release of Mostly Volcanic Carbon during the Palaeocene–Eocene Thermal Maximum’, *Nature* 548, no. 7669 (August 2017): 573–77, <https://doi.org/10.1038/nature23646>.

The CO₂ release was at least 4-5 times lower than current centennial rates of CO₂ release.¹³³ CO₂ concentrations reached 2,300ppm, with a range of 1,400ppm to 3,150ppm.¹³⁴

The Eocene Climatic Optimum was also very warm compared to today - global average temperatures were 14C warmer with a 5% to 95% range of 10°C to 18°C.¹³⁵ Temperatures remained at this level for about 6 million years. In this time, there were no substantial polar ice sheets. CO₂ concentrations were between 1,150ppm and 2,500ppm.¹³⁶

Even though temperatures were so high during the early Eocene, and even though warming was rapid on geological timescales, there was no mass extinction. In fact, the Eocene was generally a time of ecological flourishing: the name “Eocene” comes from the Ancient Greek εἶός, (“dawn”), and kainós, (“new”) and refers to the “dawn” of modern ('new') fauna that appeared during the epoch.¹³⁷

Willis and MacDonald (2011) note:

“Despite evidence for large-scale biotic turnover, little evidence suggests large-scale global plant extinction during this interval of enhanced warmth”¹³⁸

There was also a sharp increase in ordinal diversity in this period.

¹³³ “Estimated multi-millennial rates of CO₂ accumulation during this event range from 0.3–1.5 PgC yr⁻¹ 42 (Gingerich, 2019); were at least 4-5 times 43 lower than current centennial rates (Section 5.3.1.1)” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.2.3.1.

¹³⁴ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Table 2.1.

¹³⁵ “These include 7 estimates of 7°C–18°C (90% range; Inglis et al., 2020) and 12°C–18°C (95% range; Zhu et al., 2019) 8 warmer than 1850–1900, and 10°C–16°C warmer than 1995-2014 “recent past” conditions (2 standard error 9 range; Caballero & Huber, 2013). Together, they indicate that GMST was 10°C–18°C warmer during the 10 EECO compared with 1850–1900 (medium confidence).” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.3.1.1.1.

¹³⁶ “Based on boron and carbon isotope data, supported by other proxies (Hollis et al., 2019), atmospheric CO₂ during the EECO (50 Ma) was between 1150 and 2500 ppm (medium confidence), and then gradually declined over the last 50 Myr at a long-term rate of about 16 ppm Myr⁻¹ 46 (Figure 2.3).” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.2.3.1.

¹³⁷ Online Etymology Dictionary, ‘[Eocene](#)’

¹³⁸ K. J. Willis and G. M. MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’, *Annual Review of Ecology, Evolution, and Systematics* 42, no. 1 (2011): 271, <https://doi.org/10.1146/annurev-ecolsys-102209-144704>.

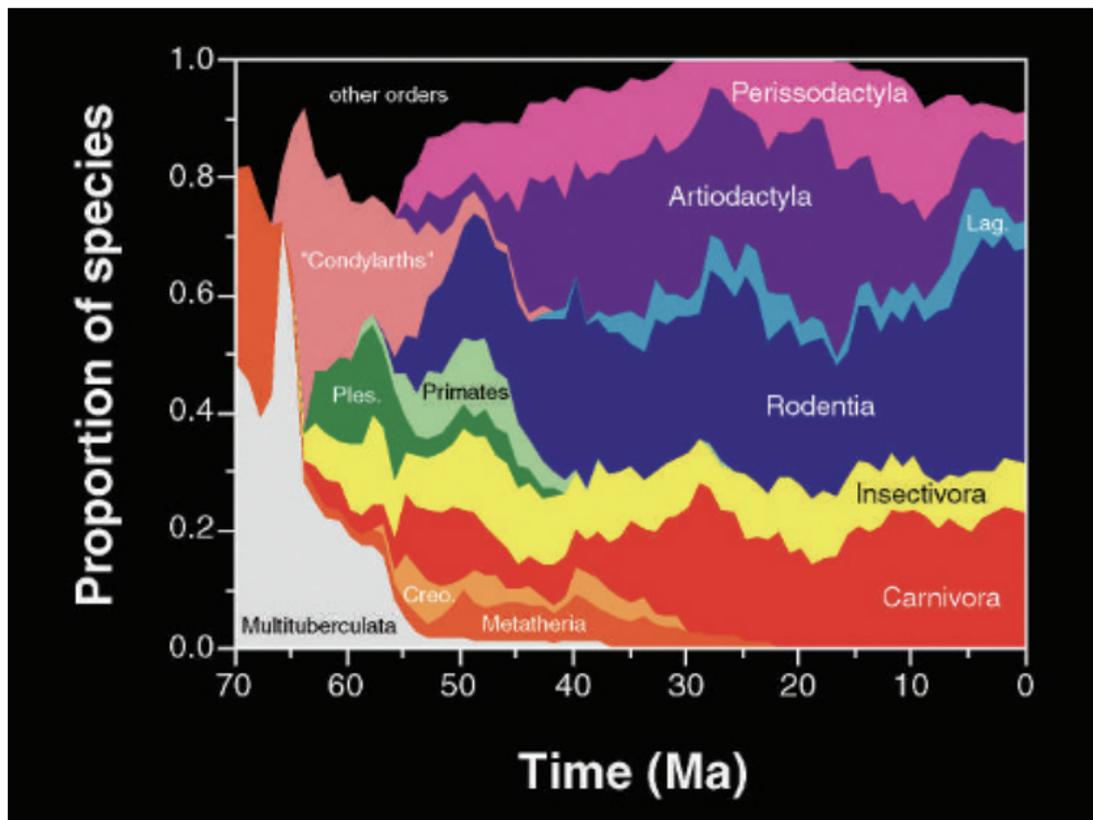


Figure 4. Cenozoic mammalian ordinal diversity plotted as a function of time. Note the sharp rises in the diversity of condylarths at 65 Ma and the secondary evolutionary bursts of Plesiadapiformes (“Ples.”) and insectivores at 62–63 Ma. By 55 Ma, all of these orders of mammals started to become far less dominant as the larger, more diversified perissodactyls, artiodactyls, carnivorans, and rodents became important parts of successive global mammalian faunas. (From Alroy, 1999; also at <http://www.nceas.ucsb.edu/~alroy/mammalorders.gif>).

Source: David E. Fastovsky and Peter M. Sheehan, ‘The Extinction of the Dinosaurs in North America’, *Gsa Today* 15, no. 3 (2005): 4–10.

Although many new fauna appeared during the Eocene, many mammals evolved transient dwarfism in order to cope with the extreme heat. This was one instance of an observation known as ‘Bergmann’s rule’, that smaller species tend to be found in warmer climates.¹³⁹ There are a number of possible explanations for this. One is that smaller animals have larger surface area per body weight in order to efficiently release heat.¹⁴⁰ Another related factor is that higher temperatures increase metabolic rate and oxygen use, and larger animals use

¹³⁹ P. B. Wignall, *The Worst of Times: How Life on Earth Survived Eighty Million Years of Extinctions* (Princeton: Princeton University Press, 2015), 70.

¹⁴⁰ Abigail R. D’Ambrosia et al., ‘Repetitive Mammalian Dwarfing during Ancient Greenhouse Warming Events’, *Science Advances* 3, no. 3 (1 March 2017): e1601430, <https://doi.org/10.1126/sciadv.1601430>.

more oxygen.¹⁴¹ During the Paleocene-Eocene Thermal Maximum, many mammals, including ancestors of modern horses, were around a third smaller in order to cope with the extreme heat.¹⁴²

The greatest negative impact of the early Eocene occurred in the ocean due to ocean acidification, thermal stress and ocean anoxia. Extinctions during the Paleocene-Eocene Thermal Maximum were limited to the benthic foraminifera, single-celled organisms that live on the seafloor.¹⁴³ Around Nigeria, there was also a massive decline in the abundance and diversity of dinoflagellates - a single-celled protist.¹⁴⁴

Overall, in the PETM, temperatures were upwards of 17°C higher than pre-industrial levels and the only species that went extinct that we know of was a single-celled marine organism, and on land it was a time of ecological flourishing, persistence and diversity. Some people paint the PETM and early Eocene as a time of great ecological trauma, which in my view is not supported by the evidence.¹⁴⁵

As Wignall writes in *The Worst of Times*:

“Part of the problem for PETM scientists who claim that there was a crisis is that they are dealing with a spectacular climate event, in particular, a story of rapid warming that was probably driven by the release of greenhouse gases, and so an extinction is anticipated (just as a pending mass extinction is predicted for modern greenhouse warming). And yet, the PETM crisis was only the faintest echo of the Pangean

¹⁴¹ P. B. Wignall, *The Worst of Times: How Life on Earth Survived Eighty Million Years of Extinctions* (Princeton: Princeton University Press, 2015), 70.

¹⁴² Abigail R. D'Ambrosia et al., 'Repetitive Mammalian Dwarfing during Ancient Greenhouse Warming Events', *Science Advances* 3, no. 3 (1 March 2017): e1601430, <https://doi.org/10.1126/sciadv.1601430>.

¹⁴³ “For years scientists considered the PETM to be the supreme example of the opposite extreme: the fastest climate shift ever known, rivaling the gloomiest projections for the future. In that light, the PETM's outcomes did not seem so bad. Aside from the unlucky foraminifera in the deep sea, all animals and plants apparently survived the heat wave—even if they had to make some serious adaptations to do so. Some organisms shrank. In particular, mammals of the PETM are smaller than both their predecessors and descendants. They evolved this way presumably because smaller bodies are better at dissipating heat than larger ones. Burrowing insects and worms, too, dwarfed.” Lee R. Kump, 'The Last Great Global Warming', *Scientific American* 305, no. 1 (2011): 56–61.

¹⁴⁴ “We attribute a massive drop in dinoflagellate abundance and diversity at peak warmth to thermal stress, showing that the base of tropical food webs is vulnerable to rapid warming.” Joost Frieling et al., 'Extreme Warmth and Heat-Stressed Plankton in the Tropics during the Paleocene-Eocene Thermal Maximum', *Science Advances* 3, no. 3 (1 March 2017): e1600891, <https://doi.org/10.1126/sciadv.1600891>.

¹⁴⁵ See for example Mark Lynas, *Our Final Warning: Six Degrees of Climate Emergency* (London: 4th Estate, 2020), 226ff. Steffen et al (2018) say

"For example, the Paleocene-Eocene Thermal Maximum (PETM) at 56 Ma BP (before present), a warming that reached 5-6°C and lasted about 100,000 years, accompanied by a rise in sea level and ocean acidification, drove the extinction of 35-50% of the deep marine benthic foraminifera and led to continent-scale changes in the distributions of terrestrial plants and animals". Will Steffen et al., 'Trajectories of the Earth System in the Anthropocene', *Proceedings of the National Academy of Sciences* 115, no. 33 (14 August 2018): SI p. 2, <https://doi.org/10.1073/pnas.1810141115>.

extinctions... What actually happened at this time is much more subtle and nuanced.”¹⁴⁶

However, there are at least two key disanalogies between the early Eocene and potential anthropogenic warming. Firstly, the rate of anthropogenic warming could be much faster than warming in the PETM. Warming of 5°C over 10,000 years is 0.05°C per century. This is very fast in geological terms, but is still two orders of magnitude slower than what we would be in for on RCP8.5, which is warming of around 4°C per century, or RCP4.5, which is around 2°C per century.

The rate of warming is likely a key determinant of the impact on ecosystems because it can affect migratory and evolutionary responses. When the climate warms, flora and fauna have to migrate to remain in their ecological niche. If the niche moves too fast due to rapid warming, then species will go extinct. Warming may also be too fast for species to adapt evolutionarily to higher temperatures. It would therefore be useful to explore cases in which past warming was as fast as future warming. One example of this is warming during the transition into the Holocene, which I will discuss below.

As Willis and MacDonald (2011) note, the second disanalogy is that

“The ecological niche apparent for many species during the Eocene may not be representative of modern-day flora. Over the past 55 million years numerous lineage splits have occurred, and these may have resulted in a loss of genetic resilience. Recent genetic work indicates that many modern species have appeared since the beginning of the Miocene [23 million years ago]”¹⁴⁷

Thus, it would be useful to explore how ecosystems coped in more recent periods.

Miocene Climatic Optimum

Period	When?	Temp vs pre-industrial	Global warming per century	CO ₂ ppm
Miocene Climatic Optimum	16 million years ago	7°C		500ppm

The Miocene epoch spanned from around 23 million to 5 million years ago. In the Miocene Climatic Optimum around 16 million years ago, temperatures were 6°C–7°C above pre-industrial levels, with the northern high-latitudes and northern temperate zones potentially 14°C and 9°C warmer.¹⁴⁸ CO₂ concentrations were 400–600ppm,¹⁴⁹ compared to

¹⁴⁶ P. B. Wignall, *The Worst of Times: How Life on Earth Survived Eighty Million Years of Extinctions* (Princeton: Princeton University Press, 2015), p. 163

¹⁴⁷ Willis and MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’, p. 274.

¹⁴⁸ “Estimates for the MCO propose global mean annual temperatures 5°C–6°C warmer than present-day with the northern high-latitudes and northern temperate zones potentially 14°C and 9°C warmer” M. Steinthorsdottir et al., ‘The Miocene: The Future of the Past’, *Paleoceanography and Paleoclimatology* 36, no. 4 (2021): 26, <https://doi.org/10.1029/2020PA004037>.

¹⁴⁹ “With higher temperatures and moderately higher pCO₂ (~400–600 ppm), the MCO has been suggested as a particularly appropriate analog for future climate scenarios, and for assessing the

business as usual CO₂ concentrations of around 550ppm. There were no ice sheets in the northern Hemisphere¹⁵⁰ and some models suggest that the Antarctic came close to being completely ice-free.¹⁵¹

Marine biota generally thrived throughout the Miocene. Steinthorsdottir et al (2021) note that:

“There were no major mass extinctions or great reorganizations and diversity among marine vertebrates appears to have been relatively high in the mid- to late-Miocene. Calcifiers such as coccolithophores, foraminifera, mollusks, and echinoderms generally thrived although a marked increase in the Mg/Ca and episodically increasing pH associated with atmospheric pCO₂ reduction may have gradually favored aragonite producers. The extended warm phase of the MCO coincided with a geographical expansion of “tropical” biota such as warm-adapted plankton and scleractinian reefs while closure of the Tethys Seaway resulted in shifts in global biodiversity hotspots.”¹⁵²

The fate of terrestrial fauna in the Miocene is particularly interesting because the epoch “has been heralded as marking the origins of “modern” terrestrial biomes as well as many of the world's biodiversity hotspots”.¹⁵³ While there was significant ecological change during the Miocene Climatic Optimum,¹⁵⁴ there was no major extinction event and there is little evidence of ecological trauma, as shown by the chart below:

predictive accuracy of numerical climate models—the same models that are used to simulate future climate.” Steinthorsdottir et al., ‘The Miocene’, 1.

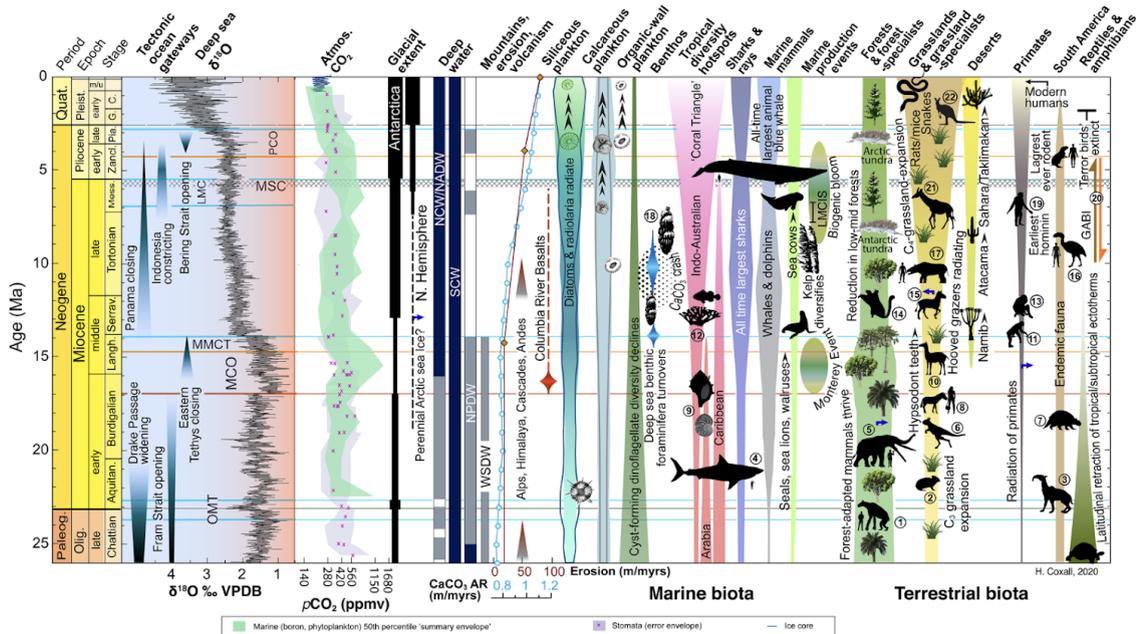
¹⁵⁰ “Continental-sized ice sheets were only present on Antarctica, but not in the northern hemisphere.” Steinthorsdottir et al., ‘The Miocene’, 1.

¹⁵¹ “These model-based estimates suggest a larger change in past Antarctic ice volume, with a minimum in ice volume close to complete Antarctic deglaciation during the MCO.” Steinthorsdottir et al., ‘The Miocene’, 33.

¹⁵² Steinthorsdottir et al., 15.

¹⁵³ Steinthorsdottir et al., 15.

¹⁵⁴ “In this section, we review the palaeobotanical record, ecosystem history, and terrestrial faunas. These records show that Miocene floras underwent the most dramatic changes of the Cenozoic, in a pattern dominated by contraction of forest biomes and replacement by grasslands, a transition that may have started already by the late Oligocene (Strömberg, 2005, 2011; Figure 4). This occurred both latitudinally and within continental interiors, reflecting the overall cooling and drying of the Earth. Although temporarily reversed during the MCO, vegetation changes were paralleled by diversification and functional evolution of mammals, resulting in our familiar flora and fauna.” Steinthorsdottir et al., 15.



Source: M. Steinthorsdottir et al., 'The Miocene: The Future of the Past', *Paleoceanography and Paleoclimatology* 36, no. 4 (2021): Fig. 1, <https://doi.org/10.1029/2020PA004037>.

Mid-Pliocene Warm Period

Period	When?	Temp vs pre-industrial	Global warming per century	Regional warming per century	CO ₂ ppm
Mid-Pliocene Warm Period	3m years ago	3°C			390ppm

In the Mid-Pliocene Warm Period around 3 million years ago, temperatures were 2.5°C–4.0°C warmer than pre-industrial levels.¹⁵⁵ CO₂ concentrations were around 390ppm.¹⁵⁶ Temperatures were much higher at the poles:

“Similar to current and predicted future climate change, the greatest warming appears to have occurred in the high latitudes, as mean annual temperatures were more than 10°C higher than present. Summer temperatures in the Beardmore region of Antarctica, for example, are estimated to have been up to 16°C warmer than present.”¹⁵⁷

¹⁵⁵ “Therefore, GMST during the MPWP is 22 estimated to have been 2.5°C–4.0°C warmer than 1850–1900 (medium confidence)” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.3.1.1.1.

¹⁵⁶ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Table 2.1.

¹⁵⁷ Willis and MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’, p. 274.

The biotic response to these temperatures is similar to the response in the early Eocene. The overwhelming response was one of global range shifts.¹⁵⁸ Willis and MacDonald comment:

“In terms of overall diversity, during this warm interval no evidence is apparent for local, regional, or global plant extinctions. Rather, where studied, evidence supports an increase in diversity. For example, on the basis of pollen-type richness, an increase in overall rainforest diversity is apparent in southeast Asia, west Africa (e.g., Morley 2000, 2007), and several sites in South America (van der Hammen & Hooghiemstra 2000, 2003). In some regions, evidence from the pollen records suggests diversity considerably higher than the present day.”¹⁵⁹

This is one piece of evidence that ecosystems closer to the present day are robust to temperatures that are 3°C higher than pre-industrial levels.

Again, there are disanalogies to the future warming. Firstly, global warming during the mid-Pliocene was much slower than potential future warming. Secondly, the spatial and temporal resolution of the Pliocene records are unable to tell us whether there were extinctions of endemic species. However, they can be addressed by examining the Last Interglacial and our current interglacial, the Holocene.

The Last Interglacial

Period	When?	Temp vs pre-industrial	Global warming per century	Regional warming per century	CO ₂ ppm
Last interglacial	127k years ago	1°C			280ppm

In the warmest millennium of the last interglacial, known as the Eemian, temperatures were 0.5 to 1.5°C above pre-industrial levels.¹⁶⁰ In high latitude regions including Greenland and Antarctica, temperatures were up to 5°C higher.¹⁶¹ CO₂ concentrations were around 280ppm¹⁶² - similar to pre-industrial levels.

Willis and MacDonald (2011) comment:

“The generally high biodiversity and remarkable resilience of all vegetation types during the Eemian are important. As far as it is possible to ascertain from the plant

¹⁵⁸ Willis and MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’, p. 274.

¹⁵⁹ Willis and MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’.

¹⁶⁰ “In summary, GMST during the warmest millennia of the 38 LIG (within the interval of around 129–125 ka) is estimated to have reached 0.5°C–1.5°C higher values than 39 the 1850–1990 reference period (medium confidence” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.3.1.1.1.

¹⁶¹ Willis and MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’, p. 276.

¹⁶² IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, Table 2.1.

fossil record, nothing went extinct despite these significantly higher temperatures and, for some species, greatly reduced distributional ranges.”¹⁶³

For all this, again, a key disanalogy may be that the *rate* of warming is especially important for ecosystem damage. We can get to grips with that issue by exploring the rapid warming that occurred in the transition into the Holocene.

Transition into the Holocene from last glacial

Period	When?	Temp vs pre-industrial	Rate of global warming per century	Rate of regional warming per century	CO ₂ ppm
Transition from last ice age to Holocene (willis and macdonald, steffensen, alley)	20,000 to 6,000 years ago			2-15°C	

Triggered by increasing summer insolation in the high latitudes of the Northern Hemisphere, the last transition between glacial and interglacial conditions commenced following the Last Glacial Maximum (20,000 years ago) of the Pleistocene and extended into the middle Holocene up to 6,000 years ago. Owing to declining summer insolation in the Northern Hemisphere in many but not all regions, a general cooling has occurred since around 5,000 years ago.¹⁶⁴ CO₂ concentrations rose from 180ppm to 280ppm.¹⁶⁵

Willis and MacDonald argue that the transition into the Holocene is particularly relevant to future warming for several reasons.¹⁶⁶

1. Similar to current warming, the glacial-interglacial transition represents a global climatic change that had repercussions from the tropics to the high latitudes.
2. The difference in global average annual temperature between glacial and interglacial conditions of 3.5°C to 5.2°C is of similar magnitude to anticipated twenty-first century warming.
3. As is the case for anticipated future warming, climatic conditions developed during the Holocene warming that had no analog during the last glacial maximum.
4. Most importantly, regional rates of warming in this period are comparable to projected future regional rates of warming.

¹⁶³ Willis and MacDonald, 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World', p. 276-277.

¹⁶⁴ Willis and MacDonald, 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World', p. 278.

¹⁶⁵ "The amount of CO₂ in the atmosphere increased from approximately 180 ppmv during the glacial to 280 ppmv during the interglacial, and this would have resulted in changes in photosynthetic activity and plant stomatal density that are expected to influence functioning, such as moisture use efficiency (Bennett & Willis 2000)." Willis and MacDonald, 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World', p. 279.

¹⁶⁶ Willis and MacDonald, 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World', p. 278-279.

- a. In Greenland temperatures may have risen by 10°C in the space of a few years, 14,700 years ago, though this may be an error in climate proxies.¹⁶⁷
- b. In the Swiss Alps and other parts of Europe, a warming of 2 to 5°C appears to have occurred in 200 years or less.
- c. In the Sierra Nevada of California, rates of warming in the late glacial may have been 4 to 5°C every 500 years around 15,000 years ago.
- d. At the higher latitudes of the Northern Hemisphere, there were increases of 5°C and more over a few decades (11,700 years ago).
- e. Data from Greenland ice cores suggest that a >10°C warming may have occurred over 20 to 60 years (13,000 to 11,000 years ago).
- f. In California, warming at the close of the Younger Dryas (11,300 years ago) may have been on the order of 3°C in less than 100 to 200 years.
- g. For the entire Southwest US, a general warming of 4°C may have occurred in less than a century (13,000 to 11,000 years ago).

During this transition, there is little evidence of plant extinctions. The predominant response was one of ecological turnover and range adjustment.¹⁶⁸

The Quaternary Megafaunal Extinction

Although no plants were lost during this period, there was an incredible loss of large land mammals, as part of the Quaternary Megafauna Extinction. Megafauna are classed as an animal weighing more than 44kg.¹⁶⁹ In North America more than 30 genera of large mammals including horses, camels, mammoths, and mastodons were lost. In South America 100% of mammals weighing >1,000 kg and 80% of mammals weighing >44 kg went extinct. In Australia, only 2 of 16 megafauna species survived.¹⁷⁰

Some of the North American megafauna - the woolly rhino, the hornless rhino, the giant ground sloth and the bear dog - are shown below:

¹⁶⁷ Dan Lunt, personal communication, 9th May 2022.

¹⁶⁸ Willis and MacDonald, 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World', p. 279.

¹⁶⁹ "But the extinctions at the end of the Pleistocene are uniquely different because they unfolded almost instantly on an evolutionary timescale and had a disproportionate bias for megafauna, a term once applied to any animal larger than a rabbit and now meaning animals with average adult body mass 44 or 45 kg (100 lbs)." G. Haynes, 'The Evidence for Human Agency in the Late Pleistocene Megafaunal Extinctions', *The Encyclopedia of the Anthropocene* 1 (2018): 219–26.

¹⁷⁰ See the Our World in Data page on [megafauna](#) for a review.

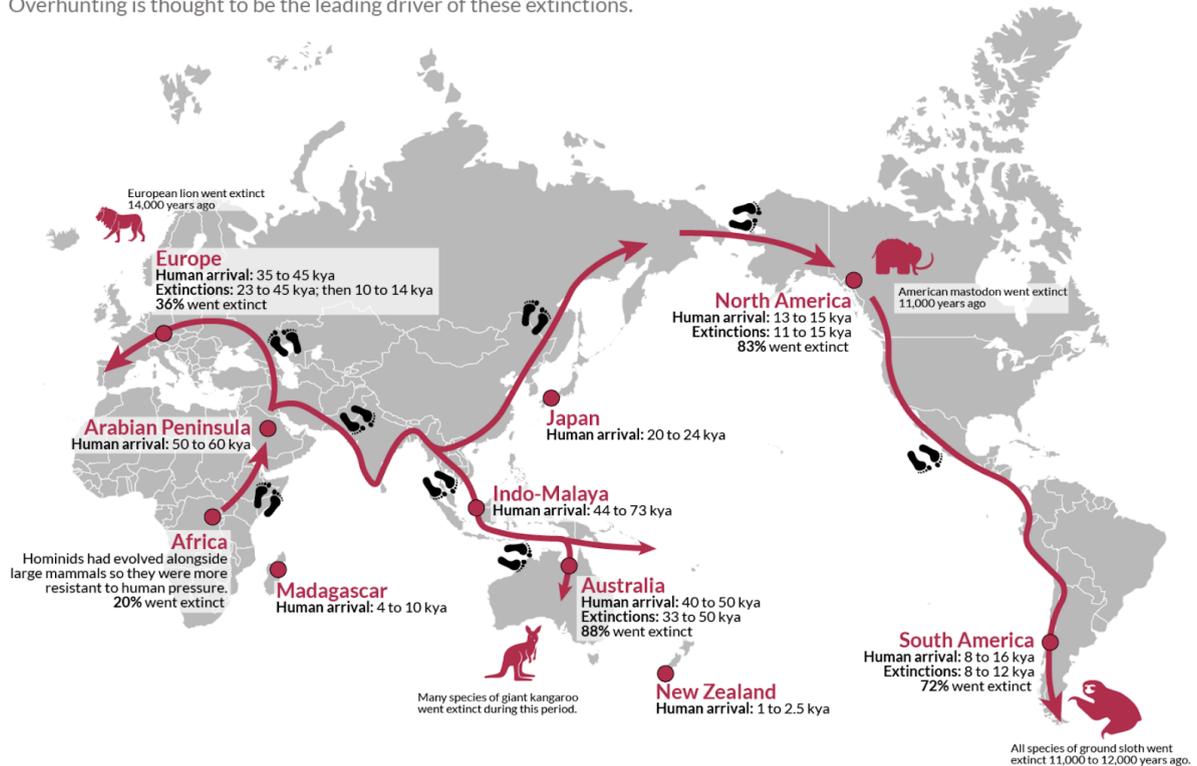


The cause of the Quaternary Megafauna Extinction is the subject of heated scholarly disagreement, with one camp arguing that human hunting killed the megafauna (aka the 'overkill hypothesis') and the other camp arguing that climate change was the culprit. I have read a lot of the literature on this topic and spoken with experts in the field and my credence in the overkill hypothesis is >90%.

There are several reasons to think that humans were the primary cause. Firstly, extinction timings on different continents closely match human arrival. At around the time that humans spread to different continents, megafauna went extinct.

Human migration and the extinction of large mammals

The Quaternary Megafauna extinction killed off more than 178 of the world's largest mammal species from 52,000 to 9,000 BC. These extinctions closely mapped human migrations across the world's continents. Overhunting is thought to be the leading driver of these extinctions.



Data Source: Andermann et al. (2020). The past and future human impact on mammalian diversity. *Science*. Images sourced from Noun Project.
OurWorldinData.org - Research and data to make progress against the world's largest problems. Licensed under CC-BY by the author Hannah Ritchie.

Proponents of overkill argue that the 'last appearance date' of a megafauna species fossil do not always overlap with the 'first appearance date' of a human fossil in an area. But this is well-explained by the patchiness of the fossil record. The last fossil we find is almost certainly not the last instance of a species and the first human fossil we find is almost certainly not the first instance of a human in an area. This is known as the 'Signor-Lipps Effect'. One illustration of this is shown by how data presented by one leading proponent of the climate change hypothesis, David Meltzer, have changed over time.

- In 1993, Meltzer noted that only **7** of the 35 extinct genera lasted until the arrival of the Clovis hunters in North America.
- In 2003, Grayson and Meltzer made the same argument, this time noting that **15** genera lasted until Clovis arrival.¹⁷¹
- In 2020, Meltzer noted that only **18** genera lasted until Clovis arrival.¹⁷²

¹⁷¹ "In 1993, Meltzer (1993, p. 306) noted that "in recent years studies of the radiocarbon chronology have shown that of the 35 species [sic] that went extinct lasted up until Clovis times." A decade later Grayson and Meltzer (2003) acknowledged that 15 of 35 extinct genera lasted until Clovis times, more than twice as many as noted before. As I write now in 2006, the current count is 17 (Stafford et al., 2005)." Gary Haynes, 'A Review of Some Attacks on the Overkill Hypothesis, with Special Attention to Misrepresentations and Doubletalk', *Quaternary International*, *World of Elephants 2*, 169–170 (1 July 2007): 89, <https://doi.org/10.1016/j.quaint.2006.07.002>.

¹⁷² "That so few of the 38 genera appear to have been hunted may be because, so far at least, only 18 of them are known to have even survived up to the time Clovis people arrived in the Americas (12)." David J. Meltzer, 'Overkill, Glacial History, and the Extinction of North America's Ice Age Megafauna',

This illustrates that the more archeology and paleontology we do, the more fossils we will find and the more that megafauna and human appearance dates will overlap.

Secondly, the Quaternary extinctions were *extremely* size-selective: the very largest animals, like mammoths and giant sloths, were preferentially killed off. The extent of the skew towards large animals is completely unprecedented in the Cenozoic.

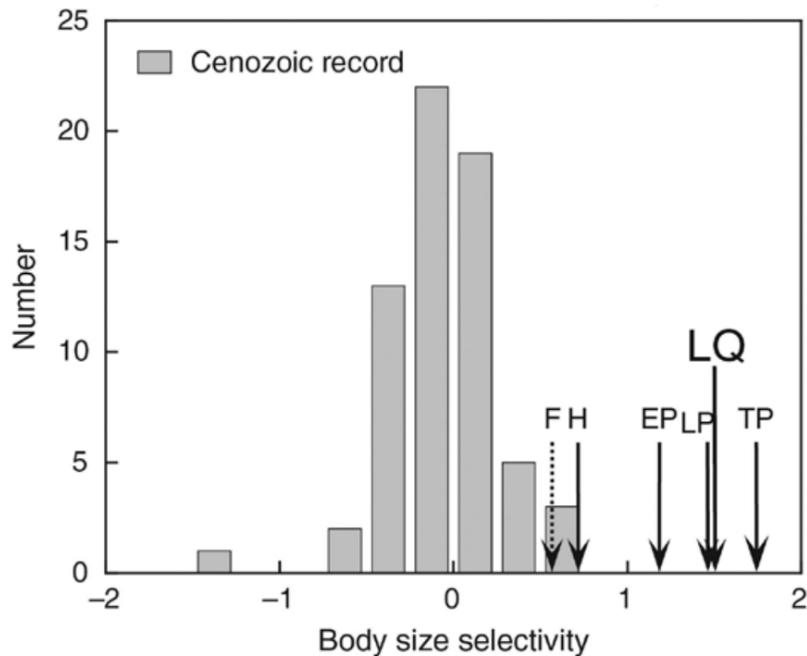


Fig. 1. Distribution of body size selectivity coefficients over the Cenozoic mammal record. All selectivity coefficients reflect change in the natural logarithm of the odds of extinction associated with a one-log₁₀-unit change in body mass. Values of zero indicate no bias, positive values indicate bias against larger size, and negative values indicate bias against smaller size. LQ, average of all late Quaternary (LP to H) extinctions; LP, late Pleistocene; EP, end Pleistocene; TP, terminal Pleistocene; H, Holocene; and F, future extinctions.

Source: Felisa A. Smith et al., 'Body Size Downgrading of Mammals over the Late Quaternary', *Science* 360, no. 6386 (20 April 2018): 310–13, <https://doi.org/10.1126/science.aao5987>.

This size selectivity is easily explained by the overkill hypothesis: larger animals offer greater rewards to human hunters. In contrast, the size selectivity is difficult to explain on the climate change hypothesis. As Wignall notes:

“The selective loss of only large animals (and those with low reproductive rates) is also not well-explained by climate change models. Under the normal ‘rules’ of extinction, highest losses generally occur among species with a relatively limited habitat range, but the Pleistocene extinctions were fundamentally different. Many of the megafaunal species inhabited a vast geographic extent: the woolly mammoth and

Proceedings of the National Academy of Sciences 117, no. 46 (17 November 2020): 28555–63, <https://doi.org/10.1073/pnas.2015032117>.

woolly rhino ranged across the whole of Eurasia and North America. Climate-driven extinction models invoke habitat change, such as the loss of tundra to advancing forests, to explain mammoth extinctions. These arguments do not account for the continuous presence of extensive tracts of all Pleistocene habitats up to the present day. In contrast, the extinctions can simply be ascribed to the observation that humans tend to hunt easy to find big animals.”¹⁷³

Proponents of the climate change hypothesis respond to this by arguing that some smaller animals also went extinct in the Pleistocene-Holocene transition.¹⁷⁴ I don't find this a compelling response to the size selectivity argument. The size selectivity argument does not require that no small animals went extinct, only that extinction rates among larger animals were hugely disproportionate.

Thirdly, the vast majority of the megafauna evolved more than a million years ago and so would have had to live through more than a dozen glacial-interglacial transitions without going extinct.¹⁷⁵ Here is one reconstruction of temperature over the last 5 million years from Hansen et al (2013):

¹⁷³ Paul B. Wignall, '6. What Happened to the Ice Age Megafauna?', in *Extinction: A Very Short Introduction*, by Paul B. Wignall (Oxford University Press, 2019), 107-108, <https://doi.org/10.1093/acrade/9780198807285.003.0006>.

¹⁷⁴ “Yet, it is important to see this episode in its broader context, for large mammal losses were not the only significant changes that took place on the Late Pleistocene landscape. Some 20 genera of birds, multiple genera of reptiles, and even a spruce tree, *Picea critchfieldii*, also went extinct at the end of the Pleistocene (Faith 2014a, Grayson 2007, Jackson & Weng 1999). Nor did the nine large mammal survivors emerge unscathed: Several, such as bighorn sheep and elk, decreased in size through the Late Pleistocene; a new species of bison arose; and there was substantial loss of genetic diversity, much of which began well before the first appearance of humans and testifies to strong selective pressures in the environment (Boulanger & Lyman 2014, Hofreiter & Barnes 2010, Orlando & Cooper 2014).” David J. Meltzer, 'Pleistocene Overkill and North American Mammalian Extinctions', *Annual Review of Anthropology* 44 (2015): 33–53.

¹⁷⁵ “Regarding the former (Table 2, rows), of the 38 genera that went extinct at the end of the Pleistocene, half of the genera (n = 19) were present throughout the entire Pleistocene, and survived multiple, previous glacial–interglacial cycles. As for the other 19 genera, six only appeared after the onset of the Rancholabrean. This puts them in North America during the MIS 7 interglacial or (using the younger age for bison arrival) during the MIS 6 glacial period. Either way, these taxa experienced just one significant glacial–interglacial cycle prior to the terminal Pleistocene. The remaining 13 genera would have had to survive at least a dozen glacial–interglacial cycles, depending on when they were first on the Irvingtonian landscape. Thus, all 38 genera experienced at least one glacial–interglacial cycle, and all survived the higher-amplitude cycles of the last 800,000 y.” Meltzer, 'Overkill, Glacial History, and the Extinction of North America's Ice Age Megafauna'.

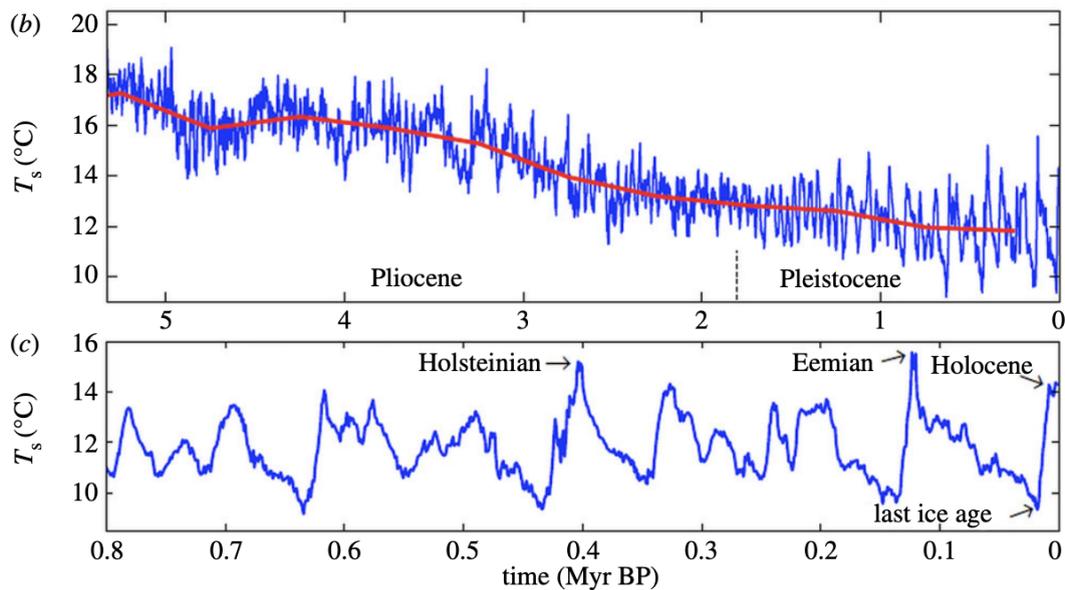


Figure 4. (a–c) Surface temperature estimate for the past 65.5 Myr, including an expanded time scale for (b) the Pliocene and Pleistocene and (c) the past 800 000 years. The red curve has a 500 kyr resolution. Data for this and other figures are available in the electronic supplementary material.

Source: James Hansen et al., ‘Climate Sensitivity, Sea Level and Atmospheric Carbon Dioxide’, *Phil. Trans. R. Soc. A* 371, no. 2001 (28 October 2013): 20120294, <https://doi.org/10.1098/rsta.2012.0294>.

As this shows, there are many comparably dramatic climatic transitions that did not cause huge species extinctions. It is only once humans made it to different continents, in the transition to the Holocene, that the megafauna started dying off. The difference between the Pleistocene-Holocene transition and the other climate transitions just does not seem large enough to cause such a discontinuous change in species extinctions.

Moreover, on the climate change hypothesis, very different kinds of climate change caused megafaunal extinctions on different continents in very different ecosystems: in some cases warming, in some cooling, in some cases aridification, in some increasing wetness. As Haynes notes:

“The species that died out in parts of [North America] that became climatically drier were the same as in the parts that became wetter; how can such opposite effects of climate changes kill every member of the same species?”¹⁷⁶

¹⁷⁶ Haynes, ‘A Review of Some Attacks on the Overkill Hypothesis, with Special Attention to Misrepresentations and Doubletalk’, 84. See also “No climatic event in North America at the end of the Pleistocene could have had such a rapid and almost simultaneous effect on so many genera with extremely variable diets (omnivory, carnivory, grazing, browsing, and mixed feeding) and living in so many different ecozones (e.g., cold northern landscapes, warmer southern regions, open grasslands, temperate forests, and semiarid environments). In fact, Late Pleistocene changes in vegetation and hydrology actually may have improved some habitats for certain now extinct megafaunal genera, such as *Mammuthus columbi* (Columbian mammoth) in temperate regions of North America, where grasslands spread as forests retreated. Late Pleistocene changes in plant communities provided adequate food for other extinct genera, such as the mixed feeder *Nothrotheriops* (Shasta ground sloth) in western North America, which successfully shifted its diet in response to different plant availability under warming temperatures, yet became extinct around the time of the first human presence. Other extinct generalist feeders in North America such as *Camelops* also should have been

It is difficult to see why these very different climatic changes happening in very different ecosystems across the world would have preferentially killed off megafauna.

One way that I think proponents of the climate change hypothesis go wrong is by focusing on extinctions on particular continents one at a time. For example, they would try to answer the question ‘what caused the extinction of the megafauna in North America?’ and effectively exclude evidence from other continents as irrelevant to this question.¹⁷⁷ But the evidence from other continents *is* relevant: we need to assess the conjunction of extinctions on multiple different continents around the time of human arrival, not each continental extinction taken individually. It is much more likely to be true that extinctions happened on multiple continents around the time of human arrival, conditional on the overkill hypothesis being true. So, taking each continent individually creates unfair bias against the overkill hypothesis.

Consider this analogy. The British serial killer Doctor Harold Shipman killed 250 of his elderly patients, mainly old women. When assessing the cause of death of each of these people, one approach would be to investigate the cause of each death in isolation and so assume that the other 249 deaths were irrelevant. If we were to do this, it would probably be rational to conclude that each of Shipman’s victims died of old age: this is by far the most likely cause of death of old age pensioners as a whole and serial killers are extremely rare. However, the deaths of Shipman’s other patients *is* relevant to a causal explanation of each death: the sheer number of deaths is much more likely conditional on the hypothesis that Shipman killed the elderly people. Ignoring the other 249 deaths ignores obviously relevant evidence. In the same way, ignoring the extinctions in South America, Europe, Asia and Australia ignores relevant evidence to the cause of the extinctions in North America. On the hypothesis that humanity was a global serial killer, we are much more likely to see extinctions happen on multiple continents at around the time of human arrival.

Fourthly, megafauna emerged relatively unscathed in Africa compared to other continents.

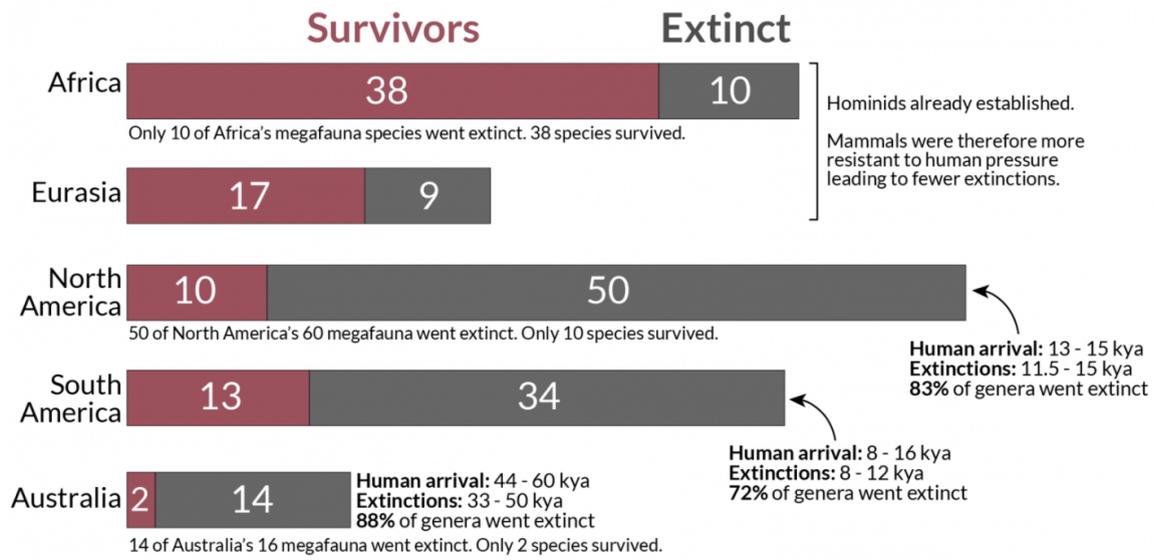
able to shift diets to accommodate available vegetation. The genus *Loxodonta* in Africa survived dramatic dietary shifts without becoming extinct, changing from predominantly grazing in the earlier Pleistocene to predominantly browsing in the Late Pleistocene and Holocene when grasslands were reduced by climate changes, demonstrating that very large animals do not necessarily become extinct when forced to shift their diets.” G. Haynes, ‘The Evidence for Human Agency in the Late Pleistocene Megafaunal Extinctions’, *The Encyclopedia of the Anthropocene* 1 (2018): 219–26.

¹⁷⁷ See for example “Although overkill has been globally applied, my focus is on Pleistocene North America south of the ice sheets” David J. Meltzer, ‘Pleistocene Overkill and North American Mammalian Extinctions’, *Annual Review of Anthropology* 44 (2015): 33–53.

Megafauna losses at the Quaternary Extinction

The Quaternary extinction event (52,000 years BC to 9,000 years BC) killed >178 species of the world's largest mammals. Humans were the primary driver of these extinctions.

Our World
in Data



Data sources: Andermann et al. (2020). The past and future human impact on mammalian diversity. *Science*. Barnkosky (2008). Megafauna biomass tradeoff as a driver of Quaternary and future extinctions. *PNAS*. OurWorldinData.org - Research and data to make progress against the world's largest problems. Licensed under CC-BY by the author Hannah Ritchie.

This is much easier to explain on the overkill hypothesis than the climate change hypothesis. Hominids evolved with megafauna in Africa, giving megafauna time to adapt to hominid hunting pressure. In other continents, this was not true and megafauna were easy prey.

It is difficult to see what the explanation for this would be on the climate change hypothesis. Furthermore, African megafauna would also have had to live through dramatic climate changes in the end-Pleistocene and yet stood a much better chance of survival.

The main response of proponents of the climate change hypothesis is to argue that there is insufficient evidence for overkill: there are too few kill sites with clear evidence of human hunting to justify the overkill hypothesis. For example, as of 2008, there were 'only' 14 documented Clovis megafauna kill sites in North America. This seems like a small number given that the Clovis would have had to have killed hundreds of millions of individuals over the course of a few thousand years.¹⁷⁸

I don't find this argument persuasive. The question we need to ask is: given the patchiness of the fossil record, how many kill sites should we expect to find? For example, one recent study of North American megafauna found only around 500 megafauna fossils out of a possible population in the hundreds of millions. It is, therefore, not surprising that there are only a handful of fossils with clear evidence of spear wounds and butchery. Surovell and Waguespack (2008) argue that the number of discovered kill sites is actually much higher than we should expect:

¹⁷⁸ "It is estimated that when Clovis hunters arrived there were hundreds of millions of these large mammals on the landscape (1). Even so, there are only 16 occurrences in which humans killed or scavenged one of these animals" Meltzer, 'Overkill, Glacial History, and the Extinction of North America's Ice Age Megafauna'.

“The United States contains more megamammal killsites than there are elephant killsites in all of Africa—a land mass that is much larger than the United States. Not only is Africa much larger, but its hominin presence extends back at least 100 times the chronostratigraphic span of the human presence in North America. Yet there are fewer than a dozen probable killsites, spanning a time range from Plio-Pleistocene to mid-Holocene”

“Obviously, given our arguments above, we should not be arguing that there are “so few” or “only about twelve” mammoth kills, but instead asking why there are so many? In over 1 million years of archeology spread over four continents, we have attempted to demonstrate that there is likely nothing that has yet to be documented archeologically that compares to Clovis in terms of the frequency of Proboscidean exploitation, with the single possible exception of the Lower Paleolithic of Iberia. Certainly, 14 sites do not seem like a very large number, but when viewed in a comparative context, it is in fact a very large number. Furthermore, the number of elephant kill sites in Clovis is truly remarkable when we consider the total number of elephant kills documented from four continents. Of the 21 sites we have identified in the Old World, only two or three have weaponry associated with carcasses. The two best cases are Lehringen and Lugovskoye, and the third case, Grobern, is questionable. Therefore, in the entire archeological records of Africa, Europe, Asia, and North America, there are a total of 16 strong cases for hunting of elephants, and 14 of these are found in Clovis. Furthermore, between of 28% and 46.7% of excavated Clovis sites that have preserved fauna are mammoth or mastodon kill sites. Why?”¹⁷⁹

Proponents of the climate change hypothesis follow one particular approach to science which puts a lot of weight on direct evidence of killing, and puts little weight on what might be deemed ‘circumstantial evidence’, which I have discussed above. Critics of the overkill hypothesis treat the question of how the megafauna died as something like a criminal trial: if you are going to claim that someone killed something, you need to actually find the fossils with spear marks. But this is not how we should form beliefs. Due to the patchiness of the fossil record, we should expect not to find much direct evidence of killing. Furthermore the ‘circumstantial evidence’ is overwhelming. From a Bayesian point of view, there is no hard distinction between the direct evidence of killing and the ‘circumstantial evidence’ I have discussed above. Circumstantial evidence can sometimes be very strong.

Another counter-argument presented by proponents of the climate change hypothesis is that humans could not possibly have killed off enough megafauna to cause extinctions. In fact, modelling evidence shows that for large, slow-to-reproduce animals, it is only necessary to kill a small percentage of the population for a species to decline rapidly.¹⁸⁰

¹⁷⁹ Todd A. Surovell and Nicole M. Waguespack, ‘How Many Elephant Kills Are 14?: Clovis Mammoth and Mastodon Kills in Context’, *Quaternary International* 191, no. 1 (2008): 82–97.

¹⁸⁰ “Quantitative predator-prey models have proven useful in studying the extinction of particular species, such as Eurasian mammoths, moas in New Zealand, or megafauna in northern Australia (Supplemental Table 4). The most comprehensive model coupled human and prey population dynamics to simulate predation on 41 large species and an undifferentiated secondary resource (plants, small game) in North America (Alroy 2001b). Hunting efficiency, the geography of invasion, and competitive interactions were varied, and all simulations assumed that hunters nonselectively

In my view, the evidence that humans were the primary driver of the Quaternary Megafaunal Extinction is very strong. Nevertheless, scholars on the topic seem to be roughly evenly divided and there are still many proponents of the climate change hypothesis. If so, one might argue that the epistemically modest thing to do would be to be agnostic and to split one's credence 50/50 between the competing hypotheses. This is an interesting test case of epistemic modesty and, I think, illustrates where the theory goes wrong. In advance of investigating the arguments on a topic, the rational thing to do might be to pick a set of experts on it and then defer to them, and one sensible set of experts to choose might be 'academics at top universities who have spent >3 years working on the question'.

However, this may not be the right way to decide what to believe *after assessing the object-level arguments* because the object-level arguments can be a good reason to give the views of some groups of experts more weight than others. While this seems epistemically immodest, as far as I know, no-one denies that it is appropriate in some cases in some cases to put less weight on the view of someone you previously believed to be an expert based on the quality of their arguments, for instance if they appeal to something widely known to be false.

One response to this is: what makes *me* better placed to judge than the proponents of the climate change hypothesis? These experts have published academic papers and books on this topic, whereas I have spent around two working weeks on it and discussed it with around ten experts. The answer to this is twofold. The first part appeals to the object-level arguments that I have outlined above. We at least need some strong object-level reasons if we are going to go against academic experts on a topic. I think that in this case, we do have some strong object-level reasons. But then, a natural response is: 'surely these experts are aware of the evidence outlined above, isn't it more likely that *you* have misunderstood something than that they are not aware of these arguments?'

This brings me to the second part of the answer: we have an *error theory* for why the critics of the overkill hypothesis go wrong. The error theory is that they do not update on 'circumstantial' evidence in a Bayesian way. Scholars do not have to take a Bayesian approach to get published in academic journals, and many scholars are not Bayesian. Although I am not an expert in megafaunal extinctions, I do have a background in epistemology, and believe I have some justification in thinking that a Bayesian approach is the correct one. If so, there is a good explanation of why many experts on the topic of megafaunal extinctions will systematically err.

took prey as encountered. Overkill occurred for a range of parameter values, although an error in the parameterization of prey r values makes it difficult to assess these results (Slaughter & Skulan 2001). In a recalculation of the best-fit trial with appropriate r values, the model correctly predicted the fate of 34 out of 41 species, with a median time to extinction of 895 years (Alroy 2001a). In that trial hunters obtained 30% of their calories from large mammals and occurred at densities of 28 people per 100 km², both within the range of values for modern hunter-gatherers. As in more generalized optimal-foraging models, the key to overkill was a relatively high human population density subsidized by smaller, faster-breeding prey. Hunting ability matters too, with greater hunting success leading to greater extinction rates, but overkill occurred even when success rates were fairly low." Anthony D. Barnosky and Emily L. Lindsey, 'Timing of Quaternary Megafaunal Extinction in South America in Relation to Human Arrival and Climate Change', *Quaternary International, Faunal Dynamics and Extinction in the Quaternary: Studies in Honor of Ernest L. Lundelius, Jr.*, 217, no. 1 (15 April 2010): 10–29, <https://doi.org/10.1016/j.quaint.2009.11.017>.

The neglect of circumstantial evidence is especially important in this case because the circumstantial evidence in favour of the overkill hypothesis is extremely strong. To recap:

- Around the time that humans arrived on multiple different continents, megafauna went extinct. This is far more likely on the overkill hypothesis than the climate change hypothesis.
- The extinctions were extremely size selective. This is far more likely on the overkill hypothesis than the climate change hypothesis: big animals are more tempting targets for hunters, but the size selectivity is nearly impossible to explain on the climate change theory.
- The climate change experienced at the end of the Pleistocene just were not different enough to preceding episodes of climate change to explain why there were such catastrophic losses of megafauna in the end-Pleistocene compared to earlier periods. This is far more likely on the overkill hypothesis than the climate change hypothesis.

Returning to our main topic, if natural climate change did not kill off the ice age megafauna, then this is further evidence that future anthropogenic climate change would not make the world radically inhospitable to humans.

3.4. Hominid flourishing and climate change

Technically, adult humans are themselves megafauna. But humans survived and flourished during the end-Pleistocene. Indeed, *homo sapiens* and our hominid ancestors have had to live through several periods of dramatic climate change. Consider again the chart from Hansen et al (2013) showing the climate since the Pliocene (5.3 million years ago).

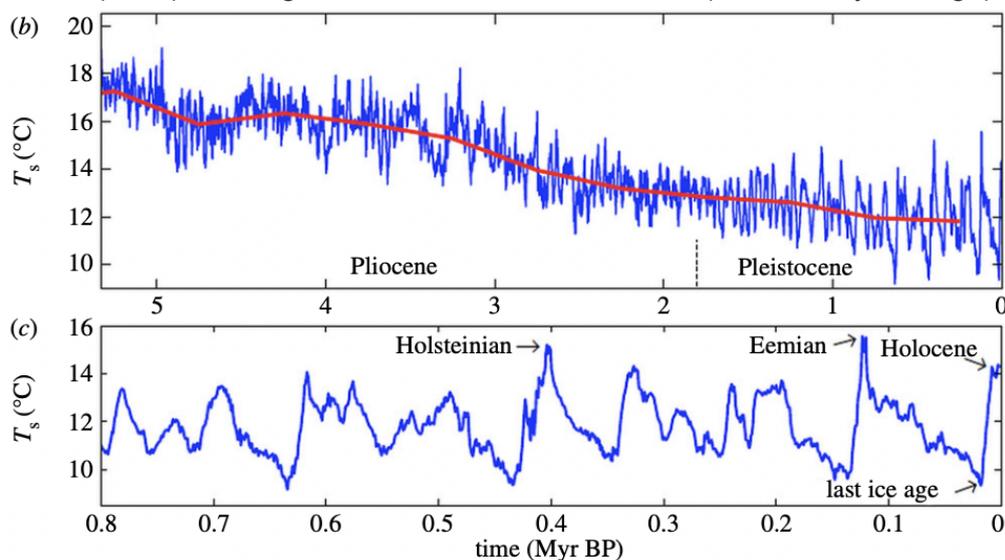


Figure 4. (a–c) Surface temperature estimate for the past 65.5 Myr, including an expanded time scale for (b) the Pliocene and Pleistocene and (c) the past 800 000 years. The red curve has a 500 kyr resolution. Data for this and other figures are available in the electronic supplementary material.

Homo sapiens evolved 300,000 years ago, and have survived through dramatic swings from glacial periods to interglacials. We survived when temperatures were 5°C lower and when

temperatures were 1°C higher. We would have had to live through the rapid regional warming in the transition from the Pleistocene into the Holocene.

Hominids evolved 6 million years ago and would have survived when temperatures were 2.5°C to 4°C warmer than pre-industrial. Indeed, at this time, the hominids would all have been in Africa, so would have had to survive temperatures that were probably 15°C warmer than modern day Britain.¹⁸¹

In many ways, our hominid ancestors were in a much more fragile situation than our own. They were much less numerous, much less technologically advanced and much less capable of a rational response to problems than people today. Thus, the survival of pre-modern hominids should provide some comfort that we will make it through 4°C of warming.

One key disanalogy between modern humans and our hominid ancestors is that we are reliant on agriculture, whereas they were hunter gatherers. I discuss the risks that climate change poses to agriculture in Chapter 5.

3.5. Why did things change after Pangea?

There is, then, discrepancy in the correlation between warming and species extinctions. Before the break-up of the Pangea supercontinent, there was a correlation, but after the break-up of Pangea, there was not. Why was this? All of the volcanic kill mechanisms I discussed above would have been in play. Why were these periods different and what does this mean for the future habitability of the planet? The answer is unclear, but one explanation points to the geology and ecology of Pangea, which made it poor at removing CO₂ from the atmosphere.

3.5.1. Pangea

299 million years ago, our planet had an unfamiliar geography. Nearly all of the world's landmasses were united into a single giant continent known as 'Pangea' that stretched from pole to pole. Pangea was surrounded by a vast ocean, even larger than the present Pacific, called Panthalassa. Pangea started to break apart 175 million years ago. Paul Wignall, a Professor of Palaeoenvironments at Leeds, has argued that Pangea was especially inhospitable to life, hence the title of his book 'The Worst of Times'. Wignall argues that, as a rule:

Massive volcanism + Pangea = major extinction¹⁸²

But once Pangea is taken out of the equation, massive volcanism barely registers on the fossil record. For instance, 135 million years ago, an eruption went to form one of the largest large igneous provinces - the Paraná-Etendeka Province. And yet, the eruption caused neither catastrophic environmental change nor mass extinction.¹⁸³ The massive eruption of

¹⁸¹ Hominids first migrated out of Africa 1.8 million years ago.

¹⁸² Wignall, *The Worst of Times*, xvi

¹⁸³ Wignall, *The Worst of Times*, 153ff.

the North Atlantic Igneous Province is one possible cause of the Paleocene-Eocene Thermal Maximum,¹⁸⁴ but as we saw above the PETM had mild ecological effects.

Wignall argues that Pangea was so inhospitable because various carbon cycle feedbacks were not in play due to the unique geography and geology of the supercontinent. By the early Cretaceous, the world was much more efficient at removing CO₂, with the result that atmospheric CO₂ concentrations were a tenth of the Pangean level.¹⁸⁵ There are several reasons for this:¹⁸⁶

- **Rainfall weathering** - In a supercontinent, huge areas are too far away from the sea to receive much rain, which reduces the scope for removal of CO₂ by rainfall and weathering.
- **Limestone deposition** - In Pangea, limestone deposition - which sequesters CO₂ in the oceans - was at a minimum because the shelf fringe of the supercontinent is much smaller than the shelf fringe of a collection of much smaller continents.
- **The evolution of coccolithophorids** - Coccolithophorids appeared in the late Triassic. They help to sequester carbon because they use CO₂ in shell formation and then sink to the bottom of the ocean when dead, which also helps to counteract ocean acidification.
- **Terrestrial plants** - The end Permian led to a mass extinction of terrestrial plants. Without plants, the weathering feedback still occurs, but plants make it happen much more rapidly. A world without plants is therefore much more prone to rapid climatic fluctuations.

Wignall provides the following diagram of the various processes of carbon sequestration:

¹⁸⁴ Stephen M. Jones et al., 'Large Igneous Province Thermogenic Greenhouse Gas Flux Could Have Initiated Paleocene-Eocene Thermal Maximum Climate Change', *Nature Communications* 10, no. 1 (5 December 2019): 5547, <https://doi.org/10.1038/s41467-019-12957-1>.

¹⁸⁵ Wignall, *The Worst of Times*, 168.

¹⁸⁶ Wignall, *The Worst of Times*, Ch. 7.

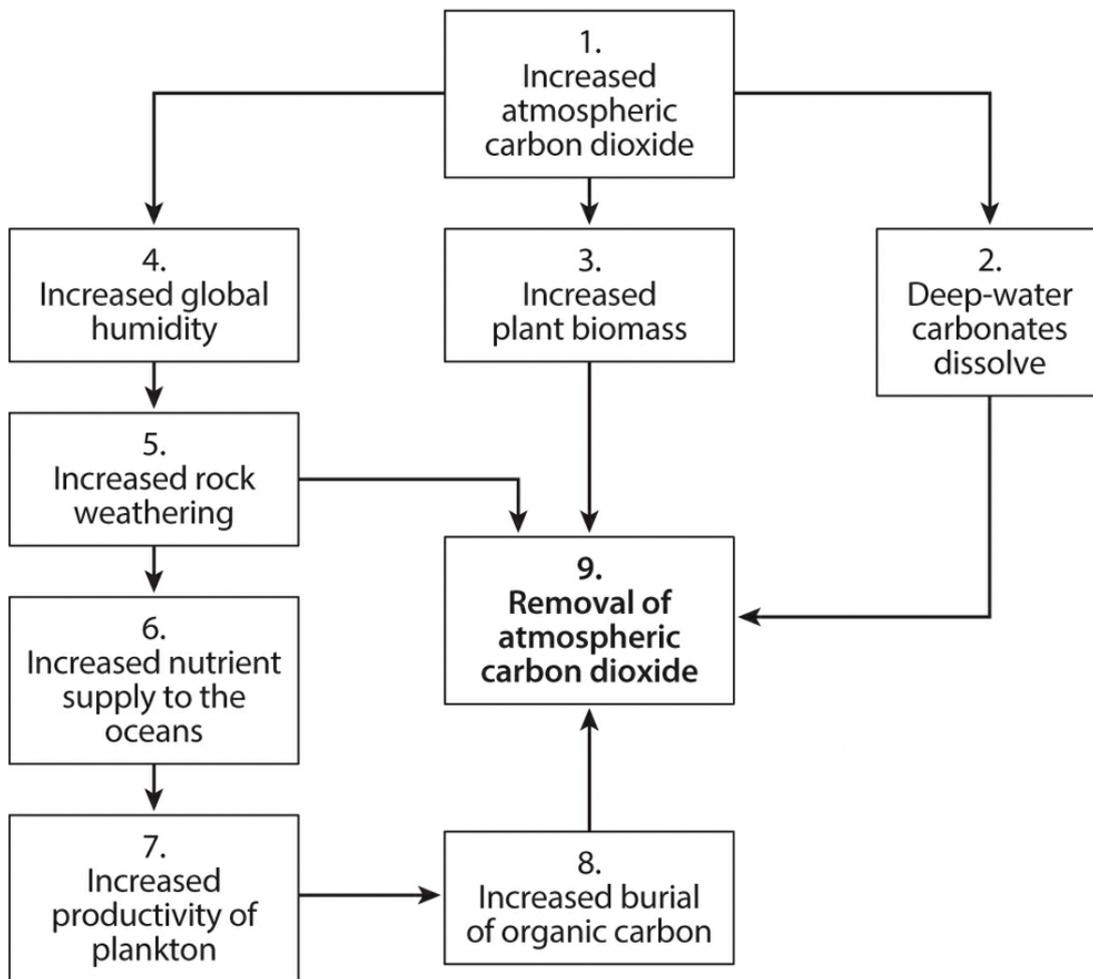


Figure 7.1. The carbon pump: how carbon dioxide is removed from the atmosphere by geological processes operating today.

Source: Wignall, *The Worst of Times*

All of this suggests that the threat the ecosystem faced in the Pangean era was peculiar to the geography and ecology of Pangea and has receded since the fracturing of the supercontinent. Mercifully, our planet may now be more resilient than it once was.

Nonetheless, our understanding of past extinction events is highly imperfect, so it may well be that we do not yet properly understand some kill mechanisms that might be relevant to future warming.

3.5.2. What about prior to Pangea?

The rate of huge extinction events prior to Pangea was also much higher from the Cambrian up until the end of the Devonian. The Carboniferous, the period between the Devonian and the Permian, was relatively peaceful. Wignall briefly notes that this may have been because of the lack of terrestrial vegetation: the first forests did not appear until toward the end of the Devonian.¹⁸⁷ I have not looked into this question in much depth.

3.6. Summarising lessons from the paleoclimate

There are no perfect paleoclimate analogues for potential future anthropogenic warming. As far as we know, there has never been a time when temperatures increased by upwards of 4°C per century on top of a baseline similar to today with something close to our current ecosystem.

But overall, I think the paleoclimate evidence provides some evidence that climate change alone will not make the planet radically inhospitable to humans. There have been cases in which our modern ecosystems and continental configuration were broadly in place and:

- Temperatures were much higher, though the rate of warming was much slower.
- Warming was comparably fast on a regional basis, albeit from a lower baseline.

In these situations, climate change was not correlated with higher rates of extinctions, nor did it lead to the extinction of hominids. Moreover, many of the previous major extinction events were driven by volcanic eruptions which not only released greenhouse gases but also toxic metals, sulphur dioxide and halogens. These other gases are the main posited cause of damage to terrestrial ecosystems but are not relevant to future anthropogenic climate change.

Our much less advanced hominid ancestors survived when temperatures were 3°C warmer than pre-industrial levels. In the Paleocene-Eocene Thermal Maximum, temperatures were 17°C warmer, but the ecological effects were surprisingly mild. There is little indication *from the paleoclimate record* that humanity would be killed off even by warming of this magnitude.

Another optimistic conclusion we can draw from all of this evidence is that, *unless something important has changed over the course of the Holocene*, future rapid warming seems unlikely to cause major species extinctions. After reviewing some of the case of warming discussed above, Willis and MacDonald (2011) conclude:

“We argue that although the mechanisms responsible for these past changes in climate were different (i.e., natural processes rather than anthropogenic), the rate and magnitude of climate change were often similar to those predicted for the next century and therefore highly relevant to understanding future biotic responses. In all intervals we examine the fossil evidence for the three most commonly predicted future biotic scenarios, namely, extirpation, migration (in the form of a permanent range shift), or adaptation. Focusing predominantly on the terrestrial plant fossil

¹⁸⁷ Wignall, *The Worst of Times*, p. 167.

record, we find little evidence for extirpation during warmer intervals; rather, range shifts, community turnover, adaptation, and sometimes an increase in diversity are observed.”¹⁸⁸

Still, there are other reasons for caution.

1. It is not yet completely clear why CO₂ release and warming were so disastrous prior to the breakup of Pangea, so there is a small chance that some unknown factor in play then might also be in play today.
2. Since there is no perfect analogue to future warming, something unexpected might happen that would affect human civilisation.
3. The world today is different in important ways to the past. Most importantly, we are reliant on agriculture for food.
4. The paleoclimate is only one line of evidence relevant to the impact of future warming. Other lines of evidence from observations and from models might paint a different picture.

¹⁸⁸ K. J. Willis and G. M. MacDonald, ‘Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World’, *Annual Review of Ecology, Evolution, and Systematics* 42, no. 1 (2011): 267–87, <https://doi.org/10.1146/annurev-ecolsys-102209-144704>.

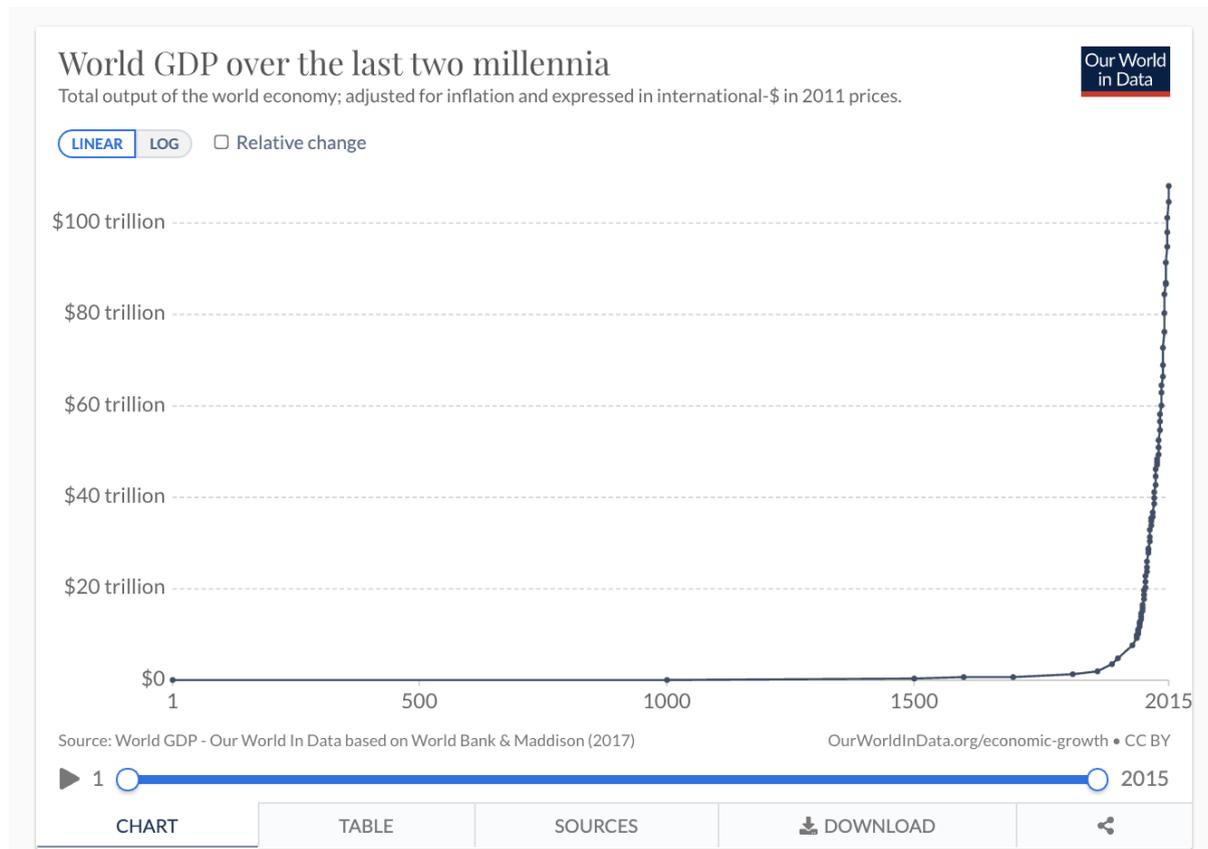
4. Economic and demographic trends

Before I discuss the effects that climate change might have on the world, it is useful to first understand the economic and demographic trends.

4.1. Economic trends

4.1.1. Global growth

I outlined the long-term economic history of the world in Chapter 1. This chart summarises the last two millennia.



Rapid growth of 1-3% has only occurred since the Industrial Revolution. Prior to that, average living standards had barely improved for thousands of years, arguably since the times of pre-agricultural hunter-gatherers.

Over the course of the 20th century, global GDP increased by 1,745%, while GDP per person increased by 350%.¹⁸⁹

It is very difficult to predict how much economic growth there will be in the future. Expert surveys suggest that income per head might increase by between 200% and 2,000% by 2100 (5% to 95% range). The shared socioeconomic pathways give a range of possibilities

¹⁸⁹ Data from [Our World in Data](https://ourworldindata.org).

for incomes in 2100, ranging from the stagnant SSP3, to the ‘middle of the road’ SSP2, to the high development SSP5.

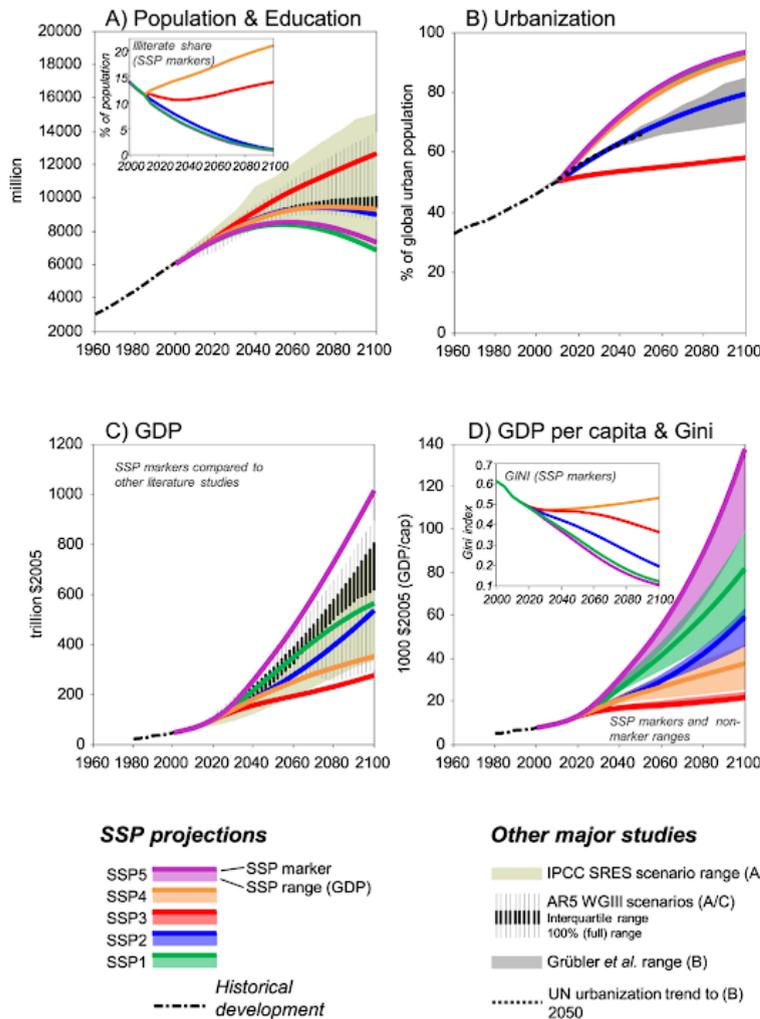


Fig. 2. Development of global population and education (A), urbanization (B), GDP (C), and GDP per capita and the Gini index (D). The inset in panel A gives the share of people without education at age of ≥ 15 years, and the inset in panel D denotes the development of the global (cross-national) Gini index. The SSPs are compared to ranges from other major studies in the literature, such as the IPCC AR5 (Clarke et al., 2014); IPCC SRES (Nakicenovic and Swart, 2000), UN, and Grubler et al. (2007). The colored areas for GDP (panel D) denote the range of alternative SSP GDP projections presented in this Special Issue (Dellink et al. (2016), Crespo Cuaresma (2016), Leimbach et al. (2016)).

Source: Keywan Riahi et al., ‘The Shared Socioeconomic Pathways and Their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview’, *Global Environmental Change* 42 (1 January 2017): 153–68, <https://doi.org/10.1016/j.gloenvcha.2016.05.009>.

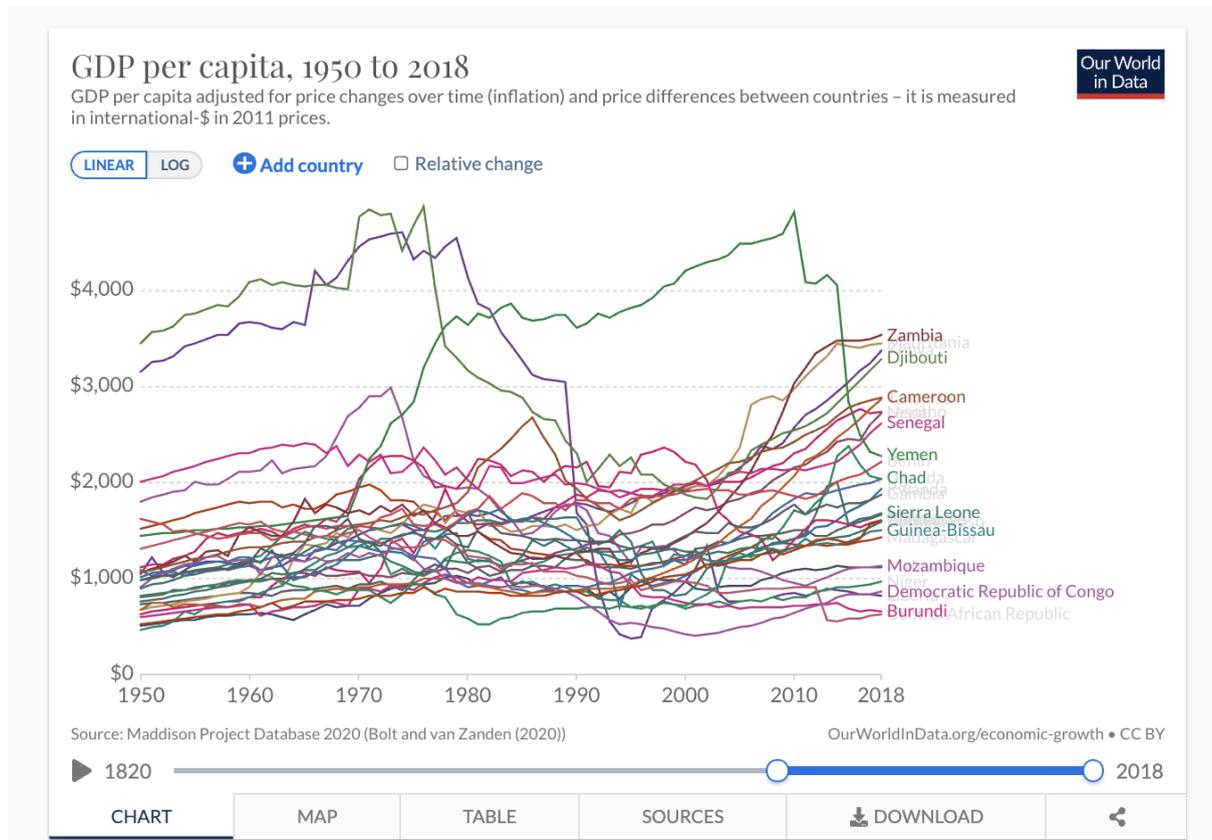
In Chapter 1, I discussed some reasons to think that this range of possibilities might be overly conservative. If advanced AI systems can allow us to automate innovation, then we could enter a new growth mode in which growth exceeds 10% per year. I would guess that the probability of this exceeds 10%.

I also mentioned that prolonged stagnation is also a real possibility. Again, I think the probability that the world economy experiences zero growth from 2100 onwards is in excess of 10%. It is also plausible that civilisation will collapse at some point in the future, so some of our descendents may be much worse off than we are.

4.1.2. Regional growth

The Industrial Revolution led to massively increased average incomes. Most of this progress was driven by growth after 1950. On all objective measures of welfare, the post-1950 era has brought more progress than [all prior human history combined](#).

However, this progress has been highly unevenly distributed. Many countries, especially in Africa, have essentially *never* experienced sustained increases in living standards. The chart below shows 29 countries, all with populations in the millions, that are still close to subsistence today and have experienced little growth over the last 70 years (though there does seem to have been some improvement since 2000).



The table below shows where different regions are headed if economic growth continues on the trend it has been on for the last 60 years.

Year	Per cap income 2019	Median growth 1960-2019	Income per head given median growth 2100
East Asia & Pacific	\$11,527	3.5%	\$186,387
Europe & Central Asia	\$24,696	1.6%	\$86,422
European Union	\$34,843	1.9%	\$159,114
Latin America & Caribbean	\$8,847	1.7%	\$33,480
Middle East & North Africa	\$8,105	0.4%	\$11,081
North America	\$63,182	2.0%	\$313,284
South Asia	\$1,960	3.6%	\$32,162
Sub-Saharan Africa	\$1,585	0.6%	\$2,520
World	\$11,429	1.9%	\$51,786

Source: [Our World in Data, 'Annual growth of GDP per capita, 1961 to 2020'](#)

Growth rates in high-income countries have started to slow in recent years, so it seems unlikely that incomes will indeed rise to such heights this century. The historical trends suggest that many countries in Africa and Asia will remain poor. However, growth rates have increased in some Sub-Saharan African countries recently.

The Shared Socioeconomic Pathways make the following assumptions about regional growth.

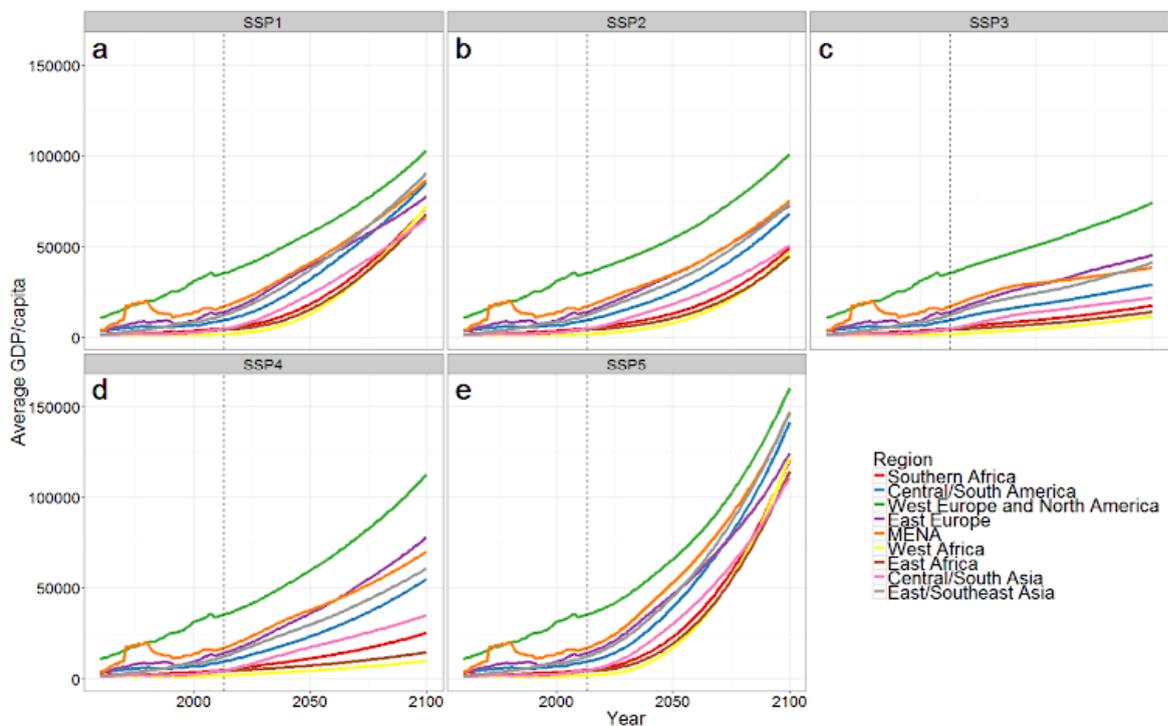


Figure S3. Country average GDP per capita (2005 USD PPP) by region and SSP.

Source: Håvard Hegre et al., 'Forecasting Civil Conflict along the Shared Socioeconomic Pathways', *Environmental Research Letters* 11, no. 5 (April 2016): <https://doi.org/10.1088/1748-9326/11/5/054002>.

Overall, the historical trends suggest that the most plausible SSP is SSP4.

Regional growth and discounting

Economists often argue that because “we” will be richer in the future, we ought to discount the future costs of climate change because people in the future will be better able to adapt to climate change than we are. In this way, paying for mitigation now can be like redistributing money from the poor (people today) to the rich (future generations).

The empirical facts I have outlined suggest that this argument is problematic.¹⁹⁰ While the *average* person seems likely to be much richer in the future, the trends suggest that a significant fraction of the world population, mainly in Sub-Saharan Africa, will not be much better off. Consequently, it is inappropriate to discount future costs using global average income growth. Rather, we need to discount future costs on the basis of likely future trends in growth in different regions, which will be highly variable.

4.1.3. Growth, impacts, adaptation and course correction

One mistake it is easy to make when assessing the future impact of climate change is to fail to take account of future trends in incomes, technology and adaptive capacity when considering the impact of climate change. For example, the only Shared Socioeconomic

¹⁹⁰ Thanks to Will MacAskill for discussion of this point.

Pathway that is consistent with RCP8.5 is SSP5. RCP8.5 promises to create some very bad climate impacts by the end of the century, but they will be borne by people who are far richer: global income per person will exceed \$100,000 in all regions. This high growth future would be as different to today as today is to 1900.

SSP5 is just a scenario, but this does illustrate an important rule of thumb: high emissions scenarios are most plausible on scenarios of high economic growth. The magnitude of climate change will likely be greatest where our adaptive capacity and ability to change course is also greatest. One exception to this is a 'boom and bust' scenario, in which there is a subsequent societal collapse

Average incomes mask significant regional variation. And, as we will see in the chapters that follow, many of the countries that are set to be hit hardest by climate change are also those with the worst prospects for growth.

The climate impacts literature typically explores the impacts of different levels of warming by exploring impacts on the most likely level of warming on a particular Representative Concentration Pathway or Shared Socioeconomic Pathway-Representative Concentration Pathway combination. For instance, many papers explore the impacts on heat stress of the most likely level of warming on RCP8.5 by 2100, which is 4.4°C. However, some papers also combine these assessments with pessimistic SSPs, such as SSP3.¹⁹¹ Technically, this combination of SSP3 and RCP8.5 is not possible, so this combination of scenarios seems inadmissible.

However, it can be useful to explore such combinations of scenarios because we should care not just about the most likely level of warming on a given RCP, but also on the risk of much higher levels of warming. On the SSP3-RCP7 baseline, there is a 5% chance of more than 4.6°C of warming.¹⁹² So, a climate impacts study which explores the impacts of SSP3 and RCP8.5, while technically impossible, is really exploring SSP3 and 4.4°C, which very much is possible.

Subsequent collapse

Although average living standards are most likely to be higher in the future, there is some chance of civilisational collapse due to nuclear war or engineered pandemics. In fact, I think the chance of a disaster killing more than 10% of the world population before 2100 is upwards of 10%. The SSPs do not account for possibilities such as this.

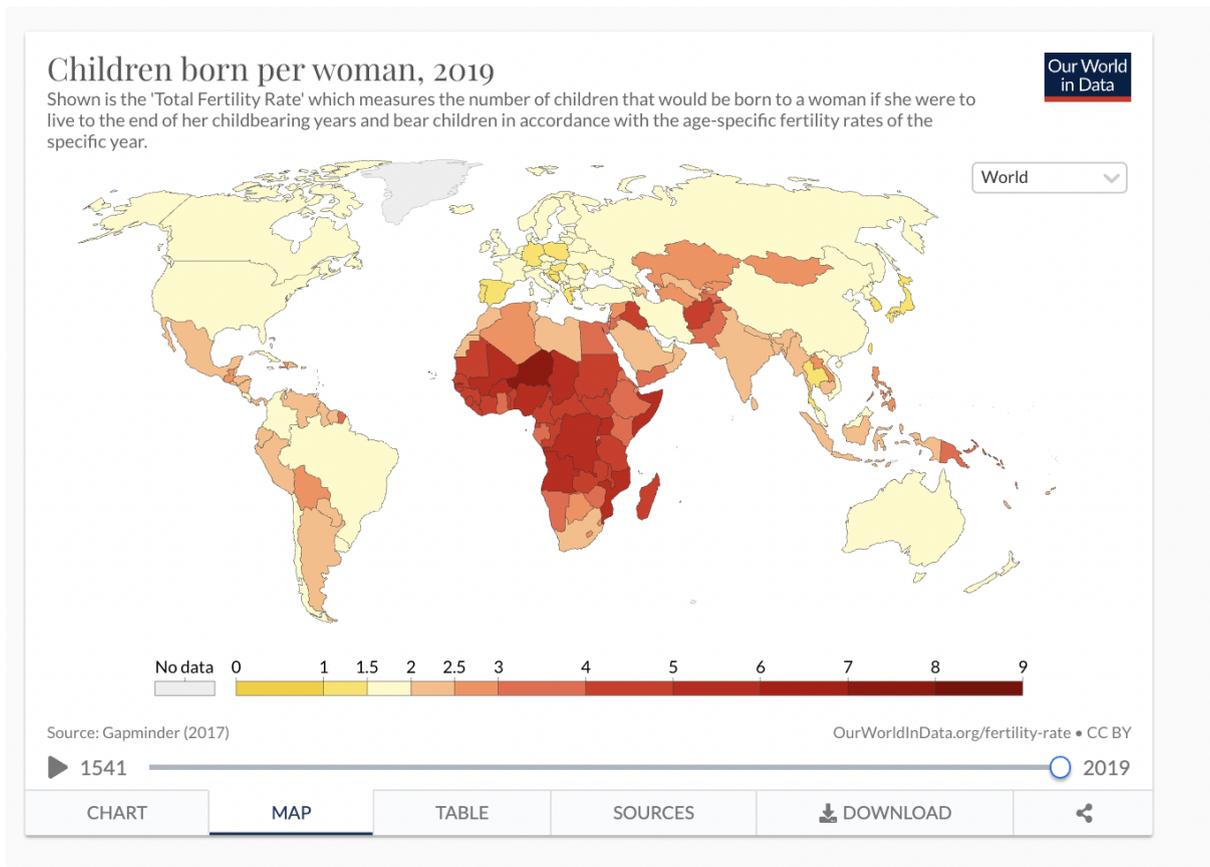
This is important because it could mean that the majority of people who have to deal with future climate change are worse-off than people today. Moreover, given how long CO₂ stays in the atmosphere, survivors of subsequent collapse will have to deal with a less hospitable climate for millennia. This is one reason that climate change could make recovery from collapse harder.

¹⁹¹ See for instance the World Bank study discussed in Chapter 11, and the Bressler et al (2021) discussed in Chapter 6.

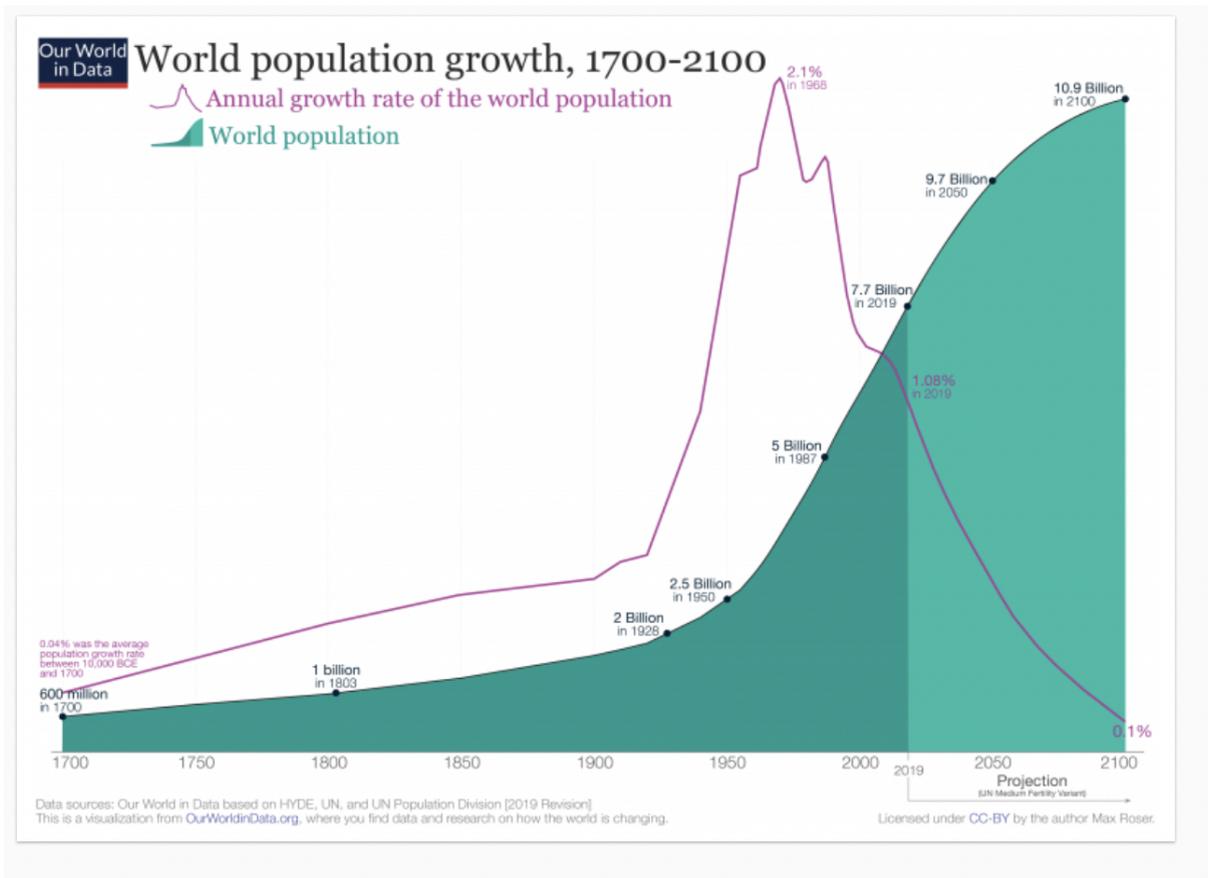
¹⁹² Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, Fig. SPM. 1.

4.2. Demographic trends

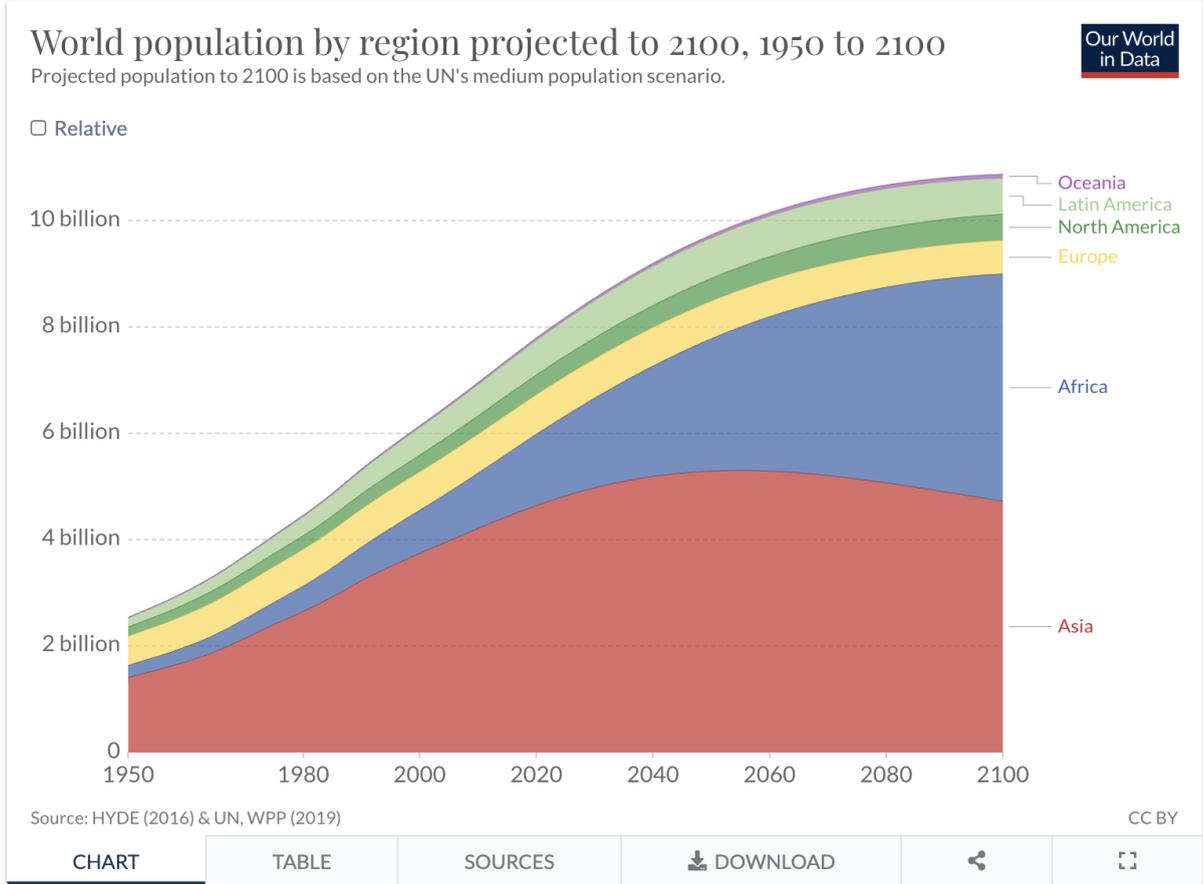
Almost everywhere outside Africa, fertility has declined close to two births per woman or below



This is a product of the [demographic transition](#): as people get richer they have fewer children. Due to rising incomes, the population growth rate peaked in 1968 and has halved since then.



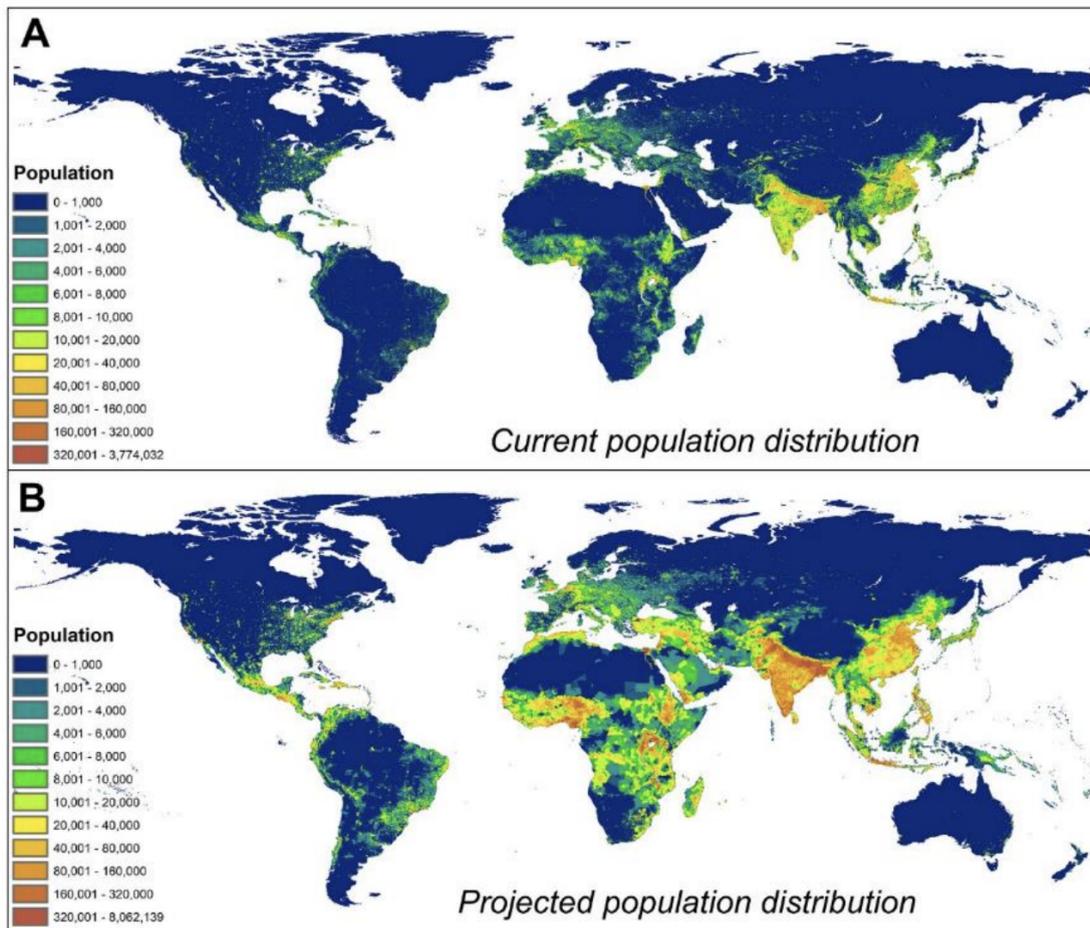
Because many African economies have been stagnant, they have not yet gone through the demographic transition. This has far-reaching implications for future global demography. The population in Africa is set to grow dramatically up to 2100 due to high fertility, and lower child mortality thanks to improved public health and medicine.



By 2100, the UN projects that 36% of the world population will be in Africa.

From the point of view of climate change, a key factor is that a lot of population growth is set to happen in the tropics and subtropics, which will be hardest hit by climate change. For example, here is the future population distribution on SSP3 (the stagnant future):

Figure S5. Global distribution of human population densities (A) as compared to the projected population distribution in 2070 following the SSP3 scenario (B). Note differences in upper bound of largest bin.



Source: Chi Xu et al., 'Future of the Human Climate Niche', *Proceedings of the National Academy of Sciences* 117, no. 21 (26 May 2020): 11350–55, <https://doi.org/10.1073/pnas.1910114117>.

5. Agriculture

In Chapter 3, we saw that early hominids survived and thrived in periods of dramatic climate change, including periods 2.5°C to 4°C warmer than pre-industrial. One disanalogy between those earlier warmer periods and potential future anthropogenic warming is that modern society is reliant on industrial agriculture, whereas our hominid ancestors would have been hunter-gatherers.

The effects of climate change on agriculture are *prima facie* some of the most potentially concerning because agriculture is heavily dependent on key climatic variables like temperature, precipitation, soil moisture, and the level of CO₂. Indeed, agriculture only started to thrive during our current warmer interglacial, the Holocene. Prior to that, *homo sapiens* subsisted by hunting and gathering. Richerson et al (2001) argue that agriculture was impossible during the cold, low CO₂ and climatically variable Pleistocene, but mandatory during the warm, high CO₂ and climatically stable Holocene.¹⁹³ As we start to move outside of our Holocene climate envelope, it makes sense to explore how agriculture might fare.

5.1. Context and trends in agriculture

5.1.1. Food production

Food production depends on:

1. Food yield - the amount of food produced per acre
2. Land area - the amount of land area used to grow food
3. Food loss - food lost between harvest and the point at which it is available to consumers.
4. Food waste - food lost at the household level.

Rice, wheat, maize and soybean produce two thirds of agricultural calories,¹⁹⁴ so I will mainly focus on them here.

Since the Industrial Revolution, the yield of the major food crops has increased enormously.

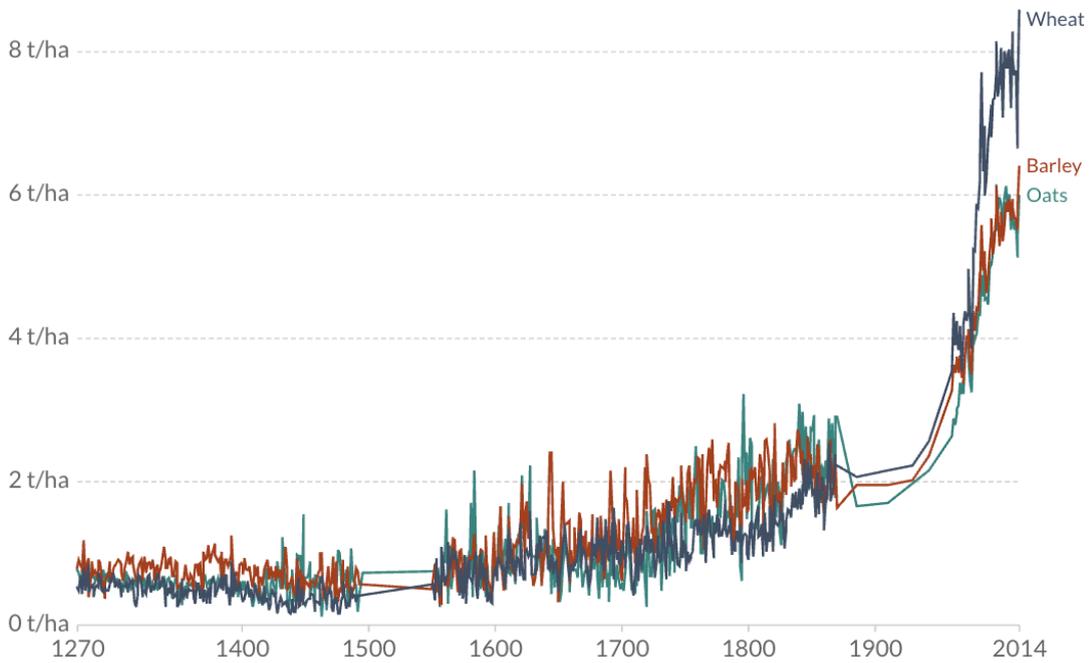
¹⁹³ Peter J. Richerson, Robert Boyd, and Robert L. Bettinger, 'Was Agriculture Impossible during the Pleistocene but Mandatory during the Holocene? A Climate Change Hypothesis', *American Antiquity* 66, no. 3 (2001): 387–411.

¹⁹⁴ "Using ~2.5 million agricultural statistics, collected for ~13,500 political units across the world, we track four key global crops—maize, rice, wheat, and soybean—that currently produce nearly two-thirds of global agricultural calories" Deepak K. Ray et al., 'Yield Trends Are Insufficient to Double Global Crop Production by 2050', *PLOS ONE* 8, no. 6 (19 June 2013): e66428, <https://doi.org/10.1371/journal.pone.0066428>.

Long-term cereal yields in the United Kingdom

Average agricultural yields in key crops in the United Kingdom, measured in tonnes per hectare.

Our World
in Data



Source: Broadberry et al. (2015) and Food and Agriculture Organization of the United Nations OurWorldInData.org/crop-yields • CC BY

▶ 1270 ○ 2014

CHART

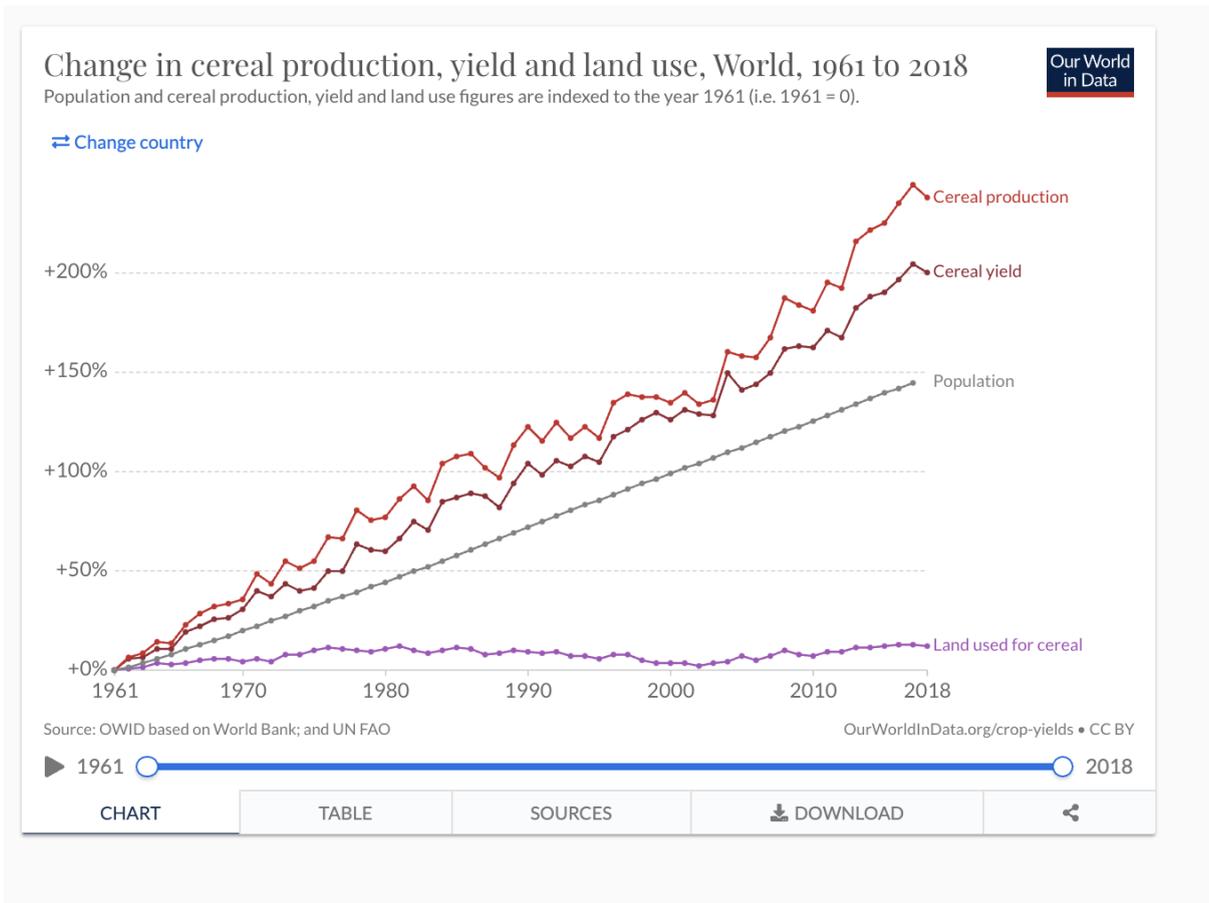
TABLE

SOURCES

DOWNLOAD



Over the last 60 years, yields and production of cereals (which includes wheat, rice, maize, barley, oats, rye, millet, sorghum, buckwheat, and mixed grains) have increased by 200% or more and have outpaced population growth.



There is a similar picture for almost all other foods, as shown [here](#). Massive increases in yield, rather than turning over land to food production, are responsible for almost all of the increase in food production. The total land area used for farmland may have [already peaked](#).

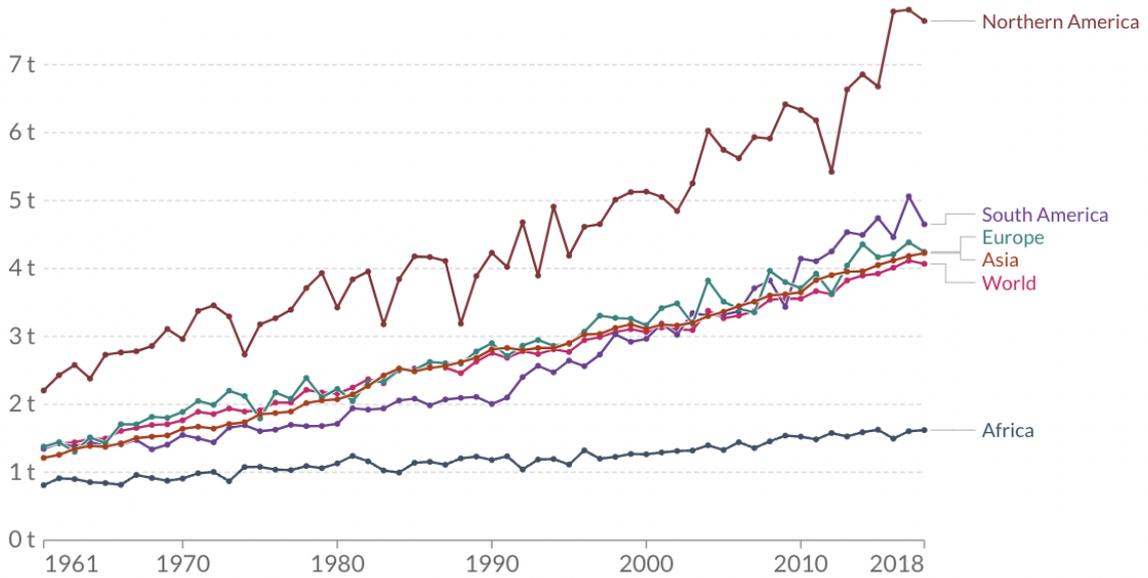
Yields have improved in all world regions, though progress has been markedly less pronounced in Africa.

Cereal yield



Cereal yields are measured in tonnes per hectare. Cereals include wheat, rice, maize, barley, oats, rye, millet, sorghum, buckwheat, and mixed grains.

[+ Add country](#)



Source: UN Food and Agriculture Organization (FAO)

OurWorldInData.org/crop-yields • CC BY

▶ 1961 2018

CHART

MAP

TABLE

SOURCES

DOWNLOAD



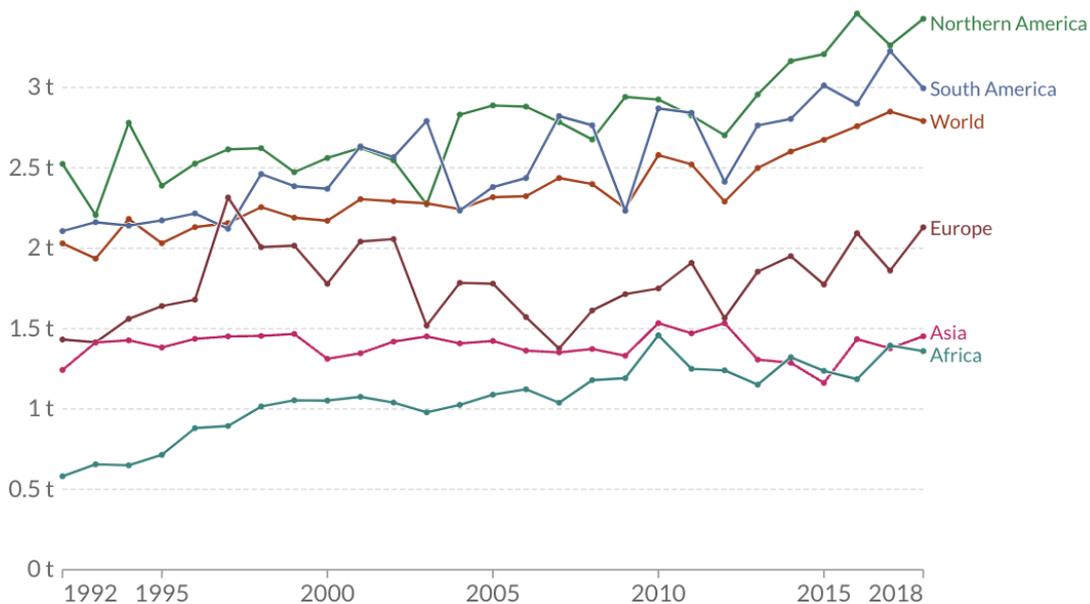
Global data for soybean yields going back to 1960 is not available, but global yields have increased by around 39% since 1990.

Soybean yields, 1992 to 2018

Average soybean yields, measured in tonnes per hectare.

Our World
in Data

+ Add country



Source: Production: Crops

OurWorldInData.org/yields-and-land-use-in-agriculture/ • CC BY

▶ 1961

2018

CHART

MAP

TABLE

SOURCES

DOWNLOAD

🔗

Much of the global improvement has been driven by the [Green Revolution](#), which involved the increased use of synthetic fertilisers and pesticides, irrigation, and higher yield crop varieties.

However, there is evidence that yields for some food crops are starting to plateau

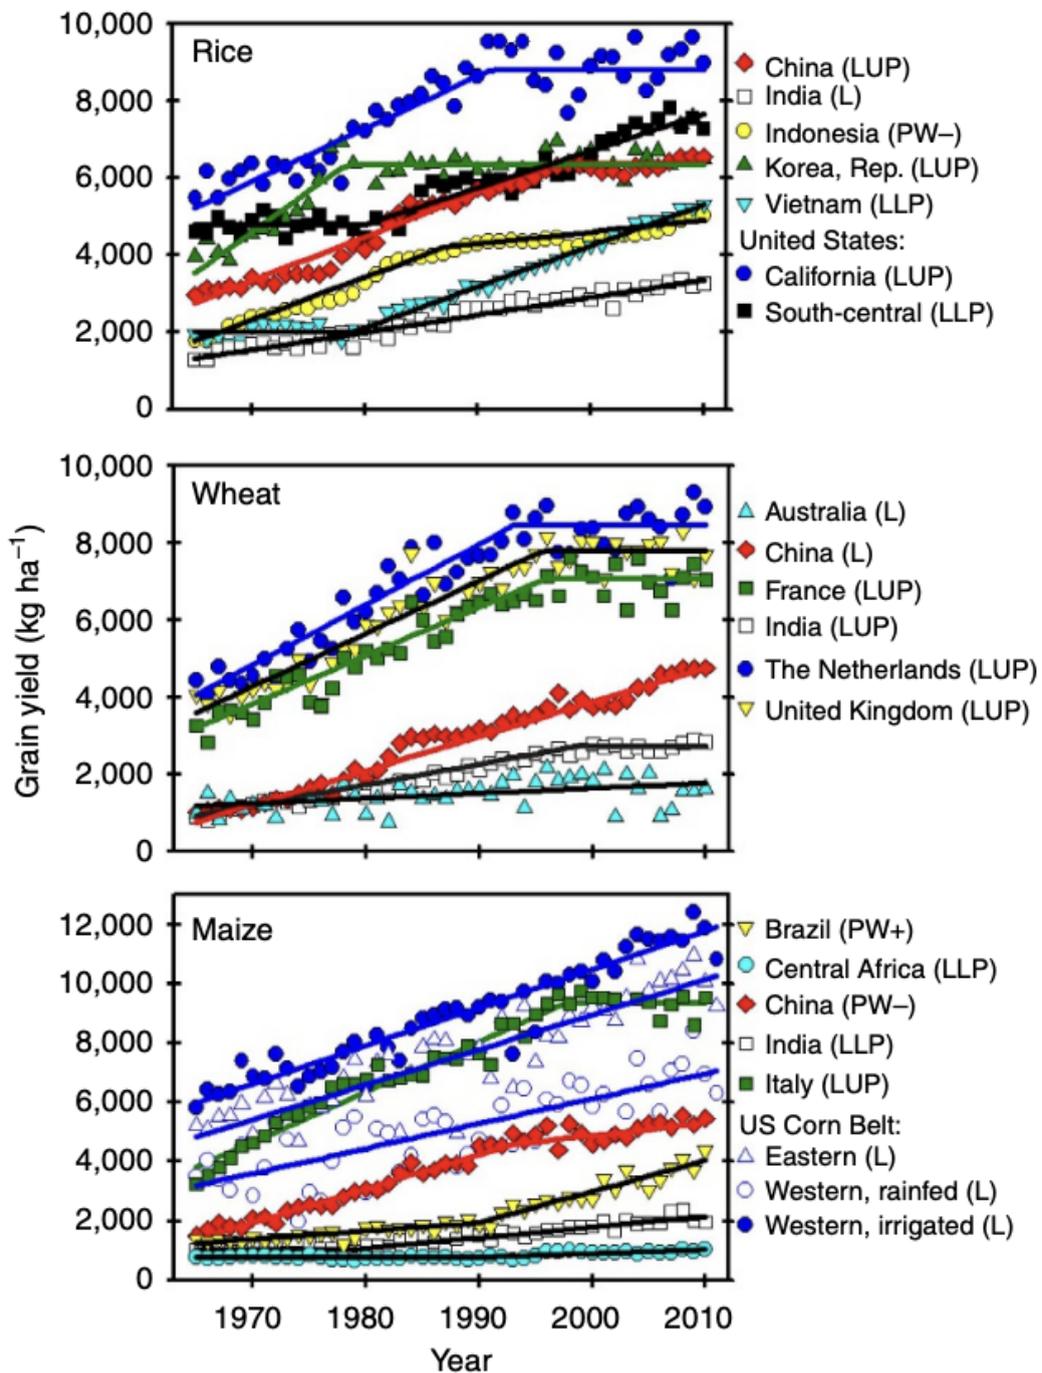


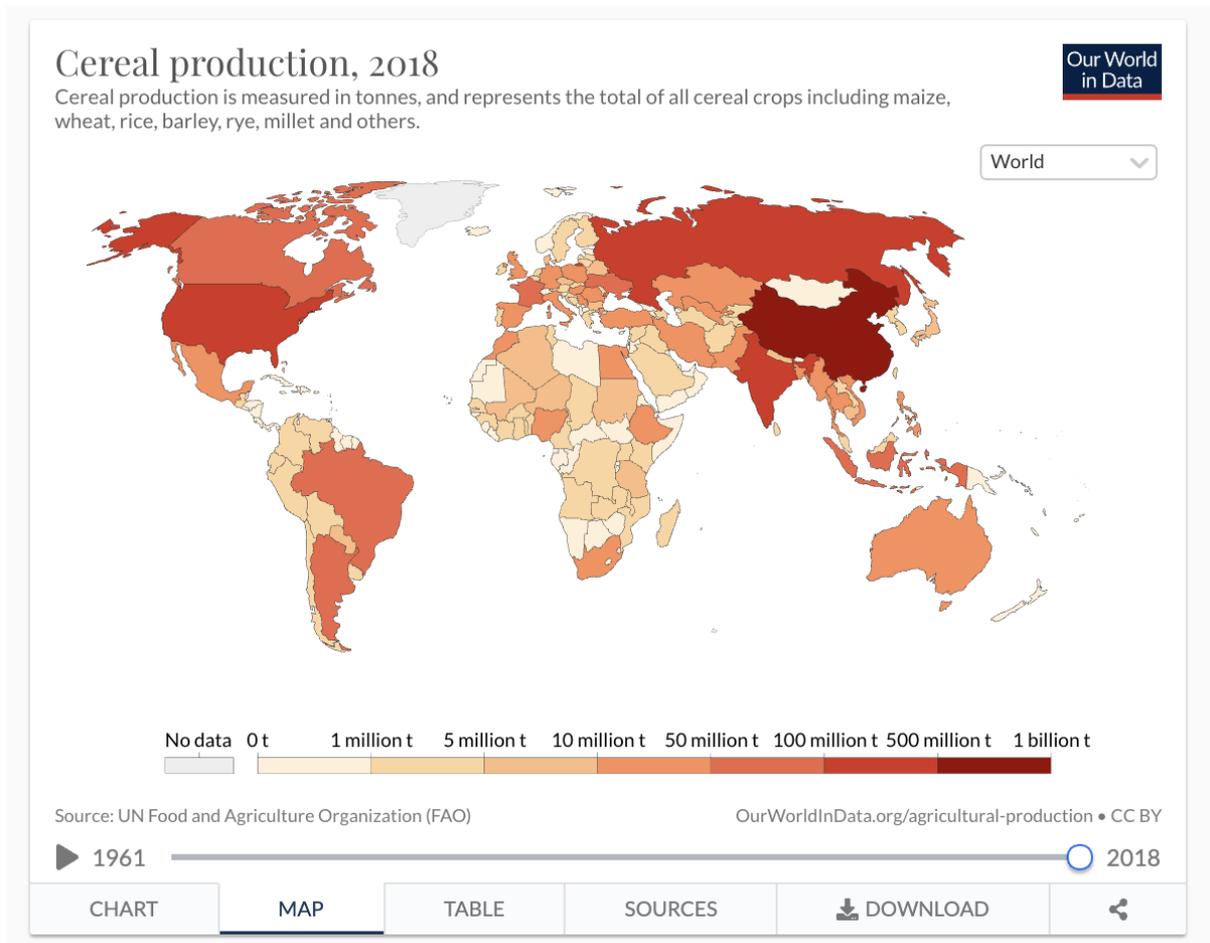
Figure 5 | Trends in grain yield of the three major cereal crops for selected regions since the start of the green revolution in the 1960s.

Fitted model for each crop-region case is indicated in parenthesis. L, linear; QP, quadratic plateau; PW, piecewise with (+) increasing or (-) decreasing rate after breakpoint year; LUP or LLP, linear with upper or lower plateau; EXP, compound exponential.

Source: Patricio Grassini, Kent M. Eskridge, and Kenneth G. Cassman, 'Distinguishing between Yield Advances and Yield Plateaus in Historical Crop Production Trends', *Nature Communications* 4, no. 1 (17 December 2013): 2918, <https://doi.org/10.1038/ncomms3918>.

Grassini et al (2013) comment that "Results from our analysis suggest that projections of crop yield trajectories based on extension of historical trends of the past five decades should be viewed with caution because these past trends were driven by rapid adoption of green revolution technologies that were largely one-time innovations."

Food production is spread fairly evenly across the globe, though Africa produces markedly less than other regions.



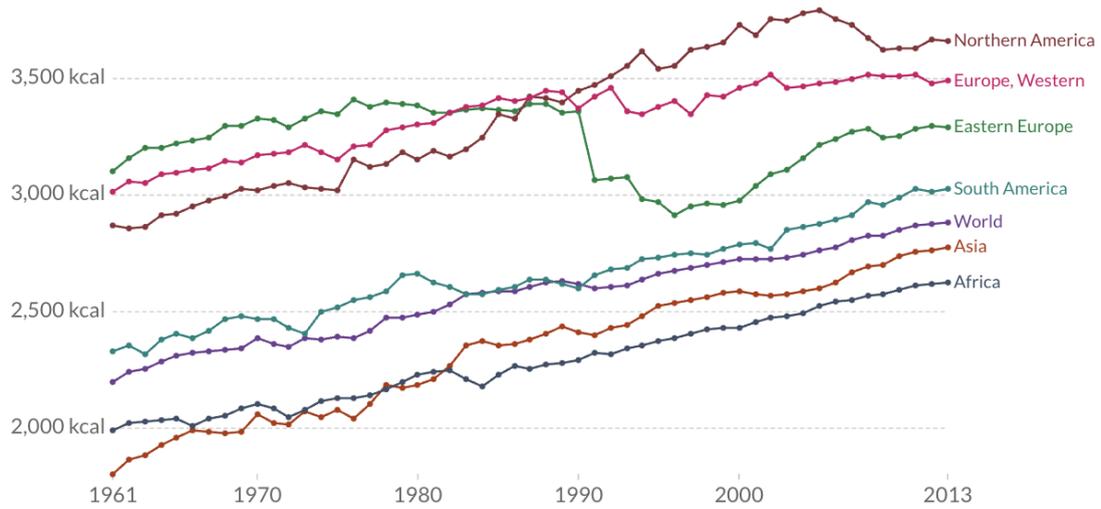
Calories per person have increased by 20% to 50% in almost all regions since 1961.

Daily supply of calories, 1961 to 2013

Caloric supply is measured in kilocalories per person per day.

Our World
in Data

+ Add region



Source: UN Food and Agriculture Organization (FAO)

Note: Data measures the food available for consumption at the household level but does not account for any food wasted or not eaten at the consumption level.

OurWorldInData.org/food-supply • CC BY

1961 2013

CHART

MAP

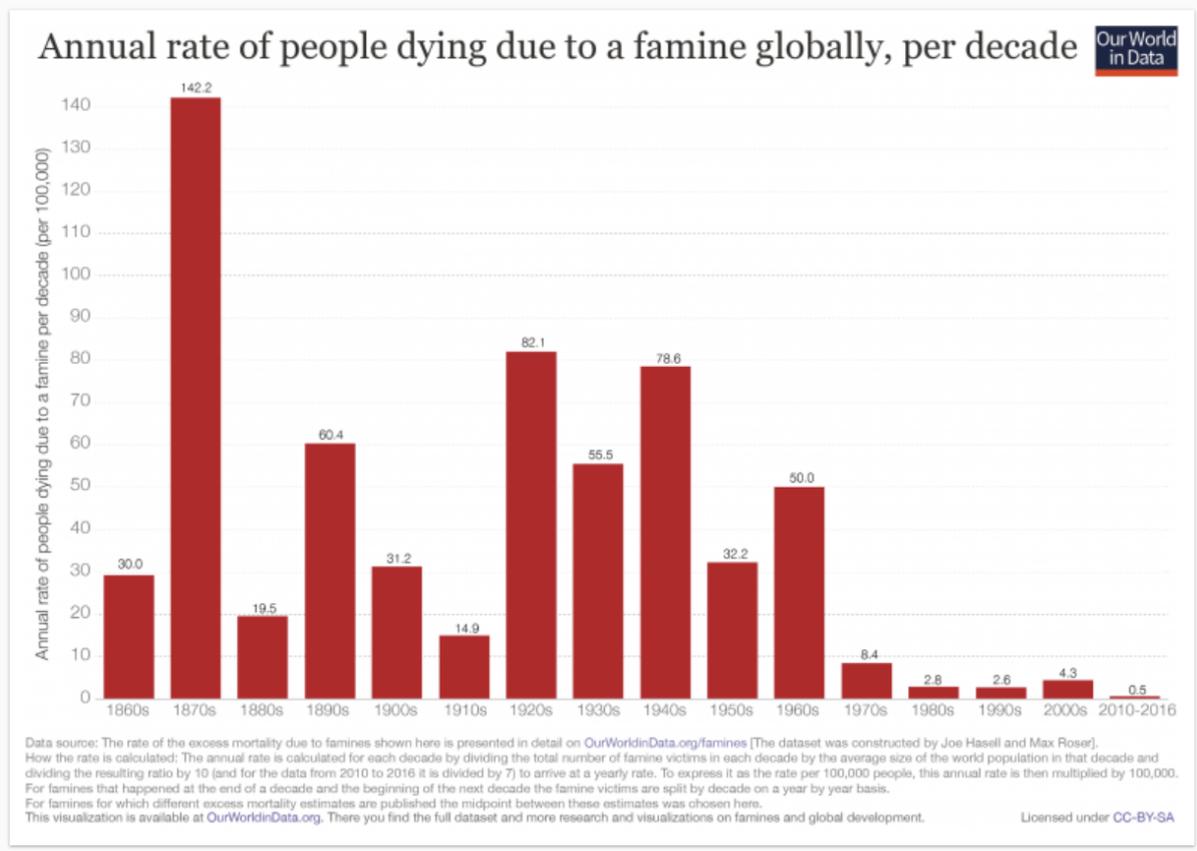
TABLE

SOURCES

DOWNLOAD



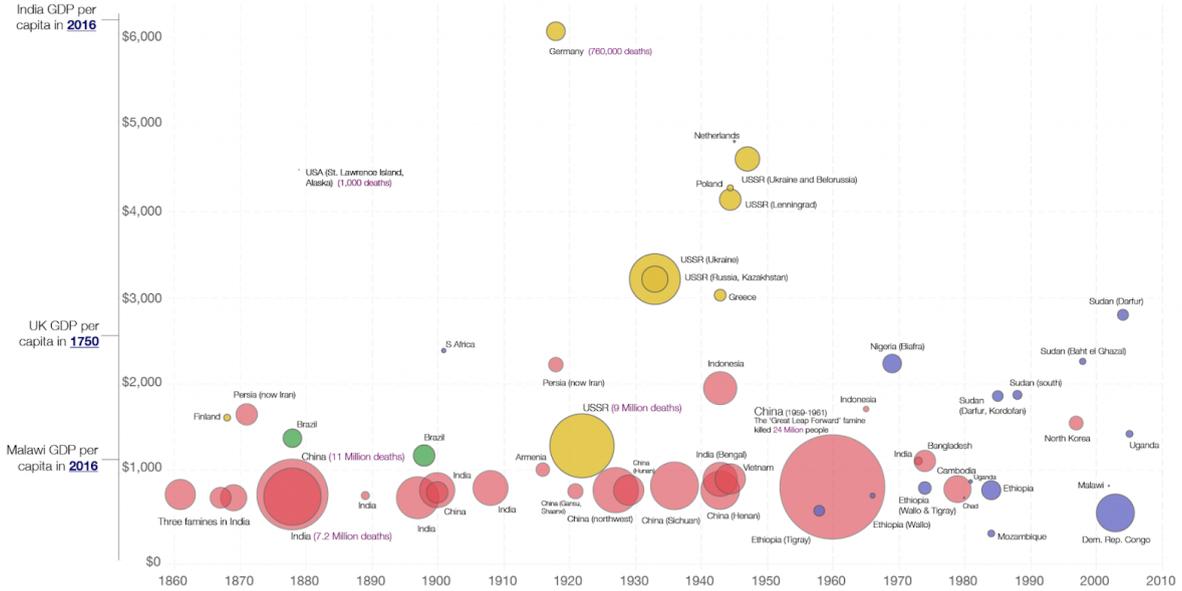
There has been a marked decline in famines over time.



Countries with income per head of more than \$5,000 tend not to suffer famines.

Famines and real GDP per capita, 1860-2016

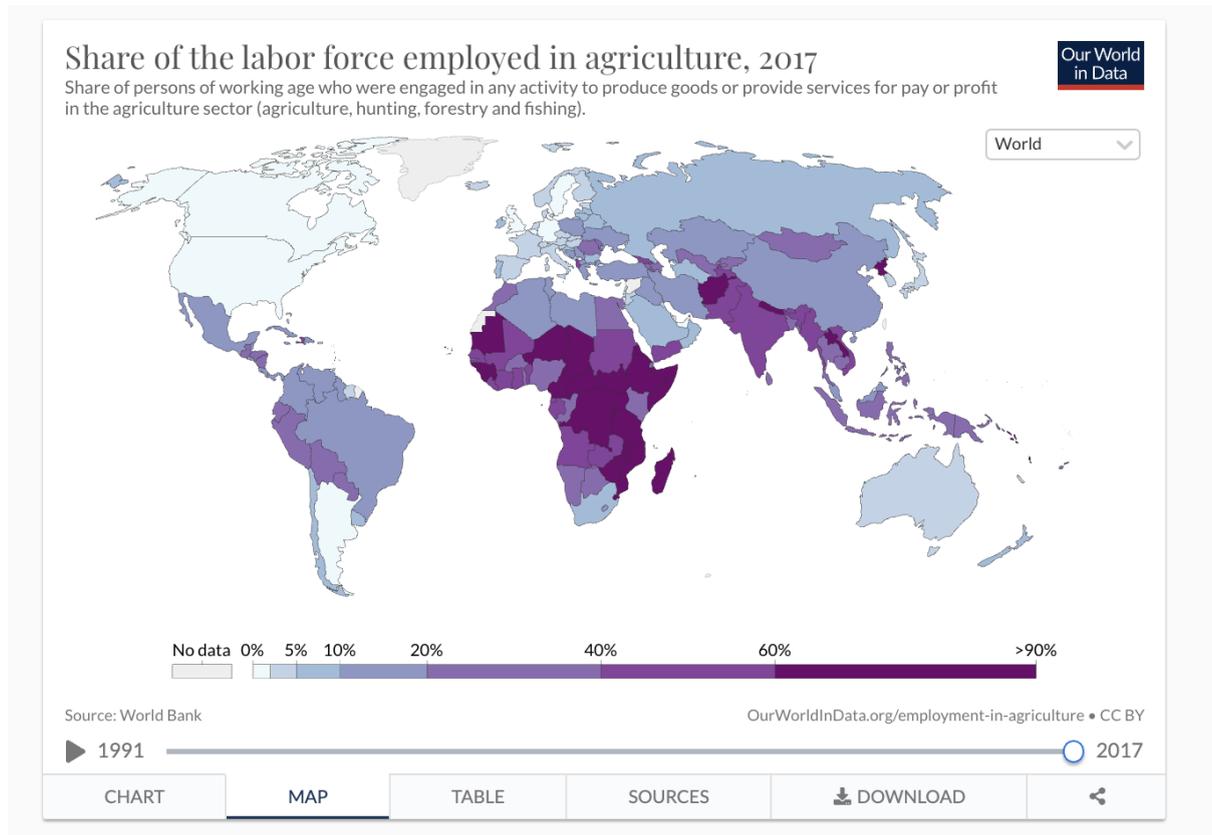
The size of the bubble represents the death count of the famine (excess mortality).
 GDP per capita is adjusted for price changes over time (inflation) and for price differences between countries (PPP adjustment)
 Detailed information on this dataset is available at OurWorldInData.org/famines.



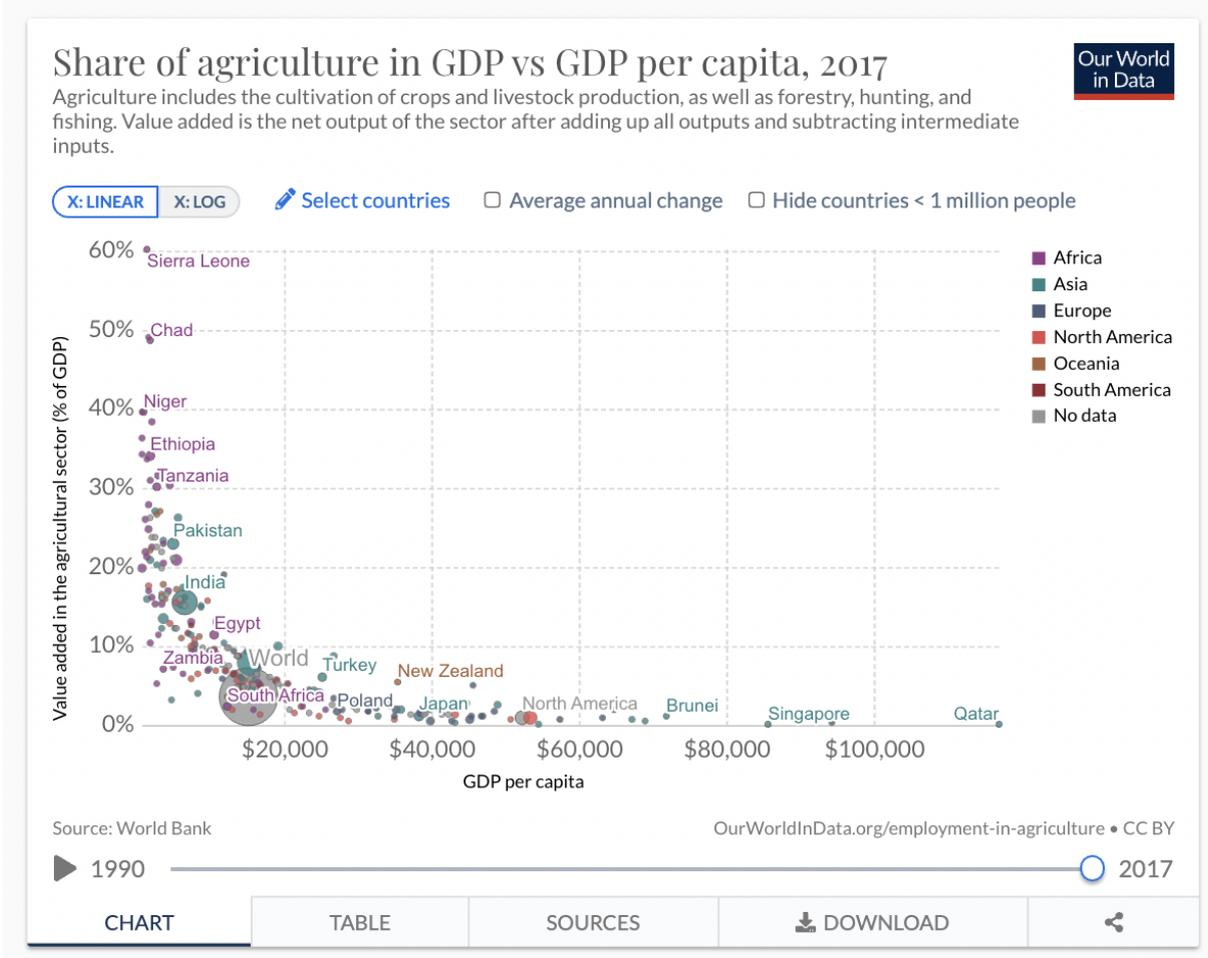
These improvements in food production have occurred in the context of around 0.8°C of warming since 1980. So far at least, global increases in yield and agricultural production have swamped the effects of climate change.

5.1.2. Agriculture and the economy

In most advanced economies, less than 10% of the labour force is employed in agriculture. In poor countries in Africa and the Middle East, more than half of the labour force works in agriculture.



Once average incomes pass \$20,000 per person, typically much less than 10% of the workforce is employed in agriculture



The [share of the global labour force employed in agriculture](#) has fallen from 43% in 1991 to 26% in 2017. Agriculture, forestry and fishing [as a share of GDP](#) has fallen from 11% in 1968 to 4.3% in 2021.

5.1.3. Historical trends in water use and irrigation

Water use

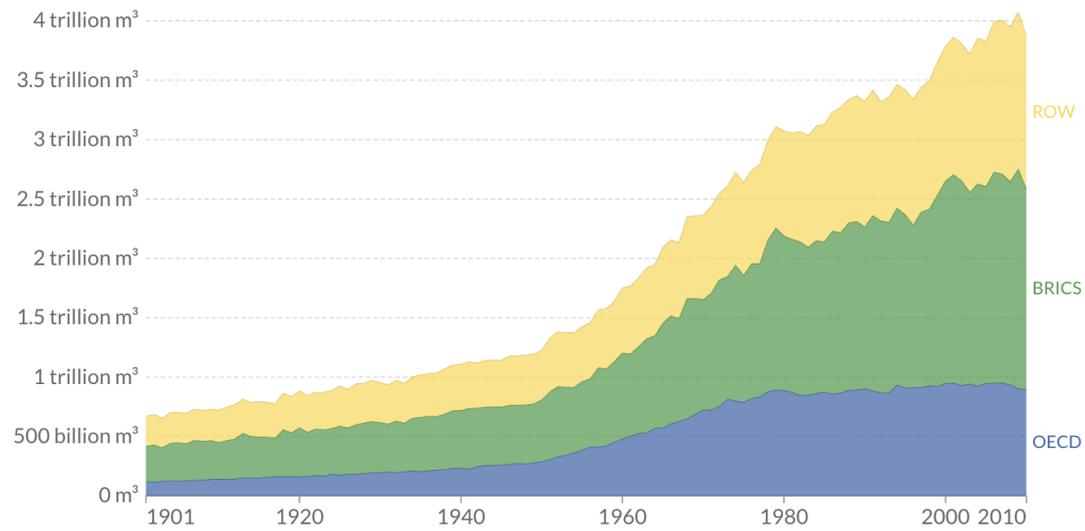
The chart below shows water consumption in different regions over the long-run.

Freshwater use by aggregated region, 1901 to 2010

Our World
in Data

Global freshwater withdrawals for agricultural, industrial and domestic uses by aggregated regional groupings. OECD members are defined as countries who were members in 2010 and their membership was carried back in time. BRICS countries are Brazil, Russia, India, China and South Africa. ROW refers to the Rest of the World, excluding OECD and BRICS countries.

Relative



Source: Global International Geosphere-Biosphere Programme (IGB)

OurWorldInData.org/water-access-resources-sanitation/ • CC BY

CHART

TABLE

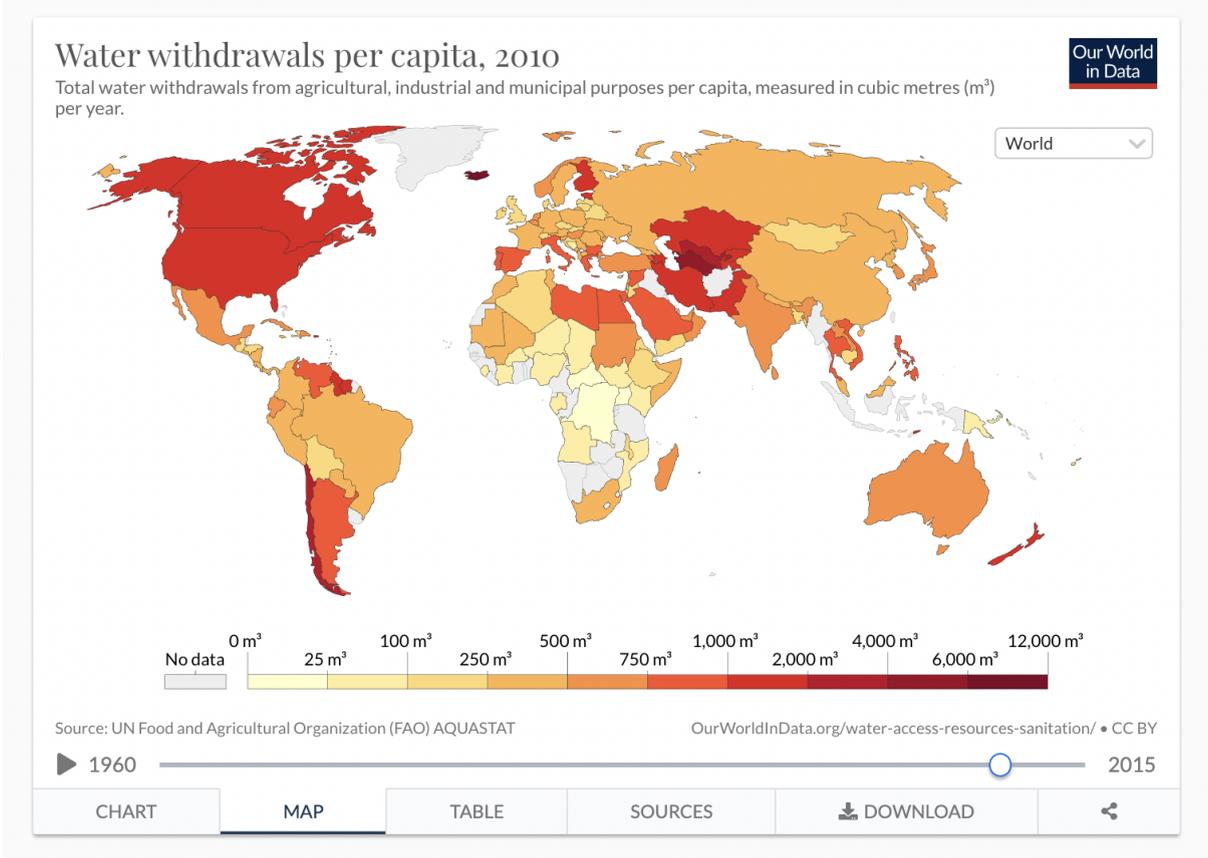
SOURCES

DOWNLOAD



Water use in rich OECD democracies has been flat since the 1980s despite economic growth, but has increased outside the OECD since 1970.

Per capita water withdrawal rates vary quite substantially across the world



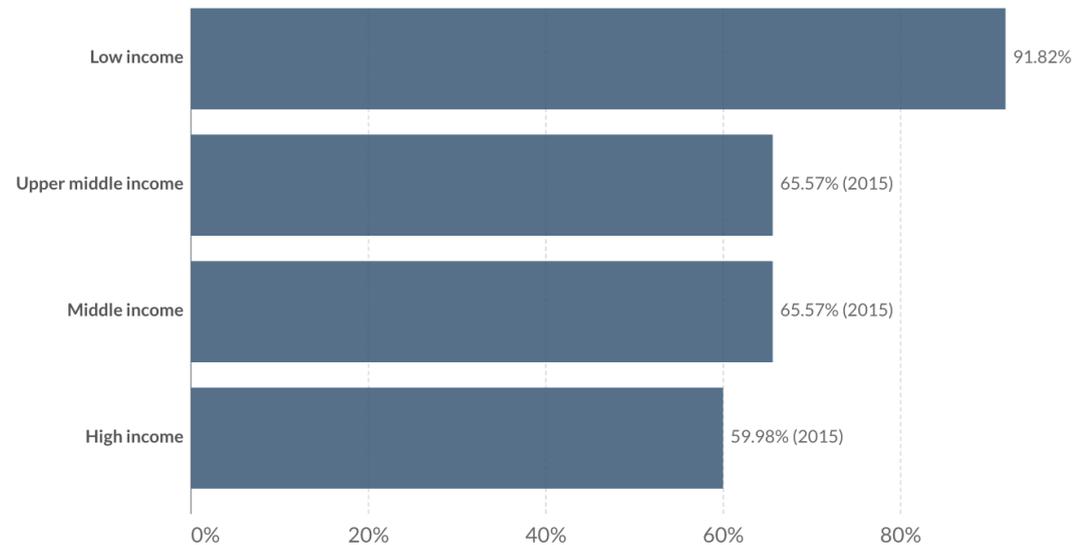
Approximately 70% of freshwater withdrawals are used for agriculture, though this varies across income groups.

Agricultural water as a share of total water withdrawals, 2016

Agricultural water withdrawals as a percentage of total water withdrawals (which is the sum of water used for agriculture, industry and domestic purposes). Agricultural water is defined as the annual quantity of self-supplied water withdrawn for irrigation, livestock and aquaculture purposes.

Our World
in Data

+ Add country



Source: World Bank

OurWorldInData.org/water-access-resources-sanitation/ • CC BY

CHART

MAP

TABLE

SOURCES

DOWNLOAD



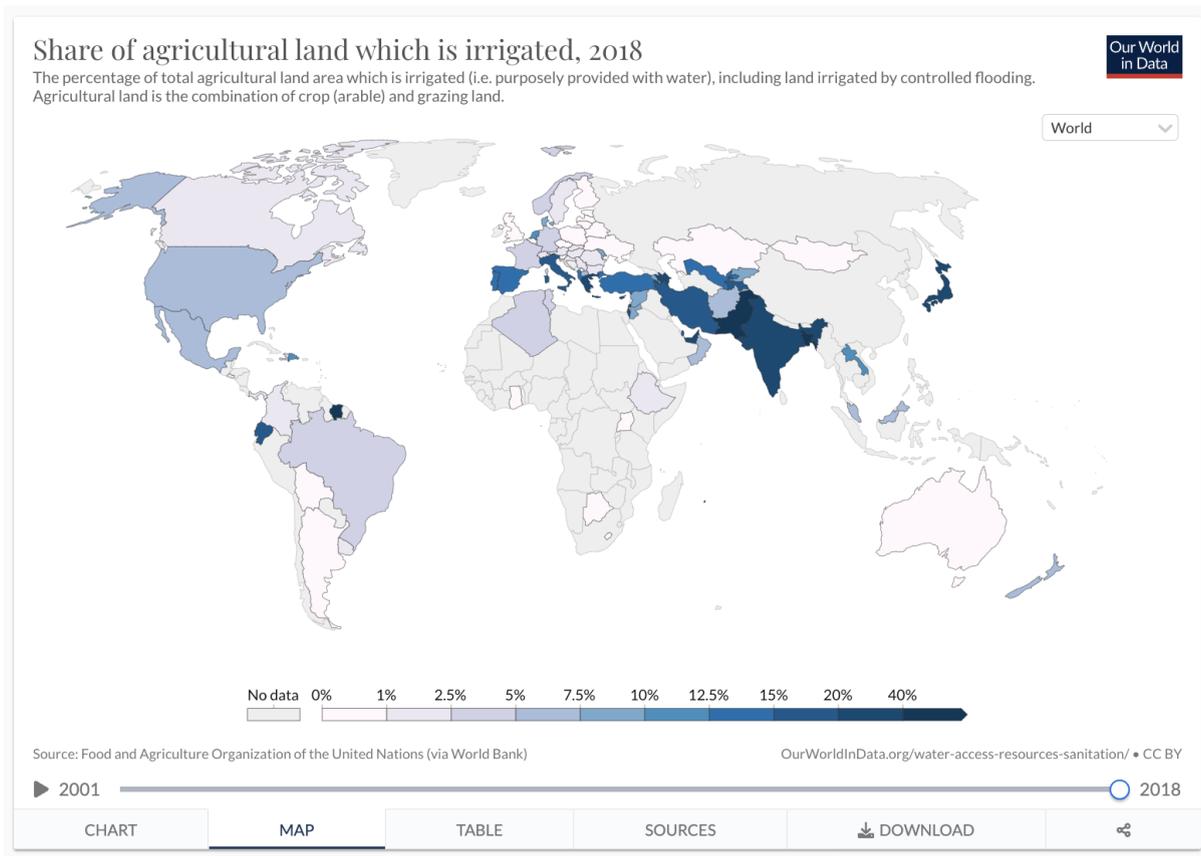
Note that water withdrawals are not necessarily environmentally harmful because water can be withdrawn but to water resources.¹⁹⁵ So, water loss is more environmentally important, though it is difficult to get global data on water loss.

Irrigation

Irrigation — the deliberate provision or controlled flooding of agricultural land with water — has been an important input factor in the observed increase of crop yields across many countries in recent decades. It has also been a strong driver in the quantity of water used for agriculture.

Rates of irrigation vary substantially across countries.

¹⁹⁵ Thanks to Linus Blomqvist for raising this point.



As this shows, even in low and middle-income countries the use of irrigation is high. In Pakistan and Bangladesh, more than half of agricultural land is irrigated, while in India around 38% is.

According to [Our World in Data](https://ourworldindata.org), low crop yields in Africa have been attributed in part to low uptake of irrigation.

5.2. By what mechanism does climate change affect food production?

Climate change can detrimentally affect the food supply in three main ways:¹⁹⁶

1. By exposing crops and livestock to increased thermal stress.¹⁹⁷

¹⁹⁶ Climate change will also influence the distribution of pests and weeds, which could also affect agricultural output. However, there is limited data on the aggregate effect of climate change on pests and weeds to date. IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, sec. 5.4.1.3.

¹⁹⁷ "Increased frequency of unusually hot nights since 1961 are also attributable to human activity in most regions (WGI AR5 Chapter 10). These events are damaging to most crops, an effect that has been observed most commonly for rice yields (Peng et al., 2004; Wassmann et al., 2009; Welch et al., 2010) as well as rice quality (Okada et al., 2011). Extremely high daytime temperatures are also damaging and occasionally lethal to crops (Porter and Gawith, 1999; Schlenker and Roberts, 2009), and trends at the global scale in annual maximum daytime temperature since 1961 have been attributed to GHG emissions (Zwiers et al., 2011). At regional and local scales, however, trends in daytime maximum are harder to attribute to GHG emissions because of the prominent role of soil

2. By increasing the risk of drought in certain regions. This effect will not be universal: in general, models suggest that climate change will make dry areas like the subtropics drier and wet areas at higher latitudes wetter.¹⁹⁸
3. By exposing agricultural workers to increased heat stress and flooding.

However, CO₂ emissions also have some beneficial effects on crops.

1. Higher temperatures extend the growing season, which is especially beneficial to countries at higher latitudes.¹⁹⁹
2. Elevated CO₂ levels speed up photosynthesis and water use efficiency due to the CO₂ fertilisation effect. The meta-analysis used in the 2013-14 IPCC reports found that each ppm increase in CO₂ increases yield by 0.06%.²⁰⁰ On RCP6, concentrations would increase by 300ppm compared to today, so yields would increase by 18%, other things equal.

5.3. What is the projected effect of climate change on agricultural production?

5.3.1. Crop yields

In this section, I will outline the findings of major reviews on the effects of climate change on crop yield. It is important to note that these studies measure impacts relative to a counterfactual without climate change, not relative to today. While the world warms there are likely to be countervailing improvements in agricultural productivity, I point I return to below.

IPCC Fifth Assessment Report

The chart below shows the results of a meta-analysis used in the IPCC's Fifth Assessment Report, on the effect of warming on food yield in different regions. The adaptations explored

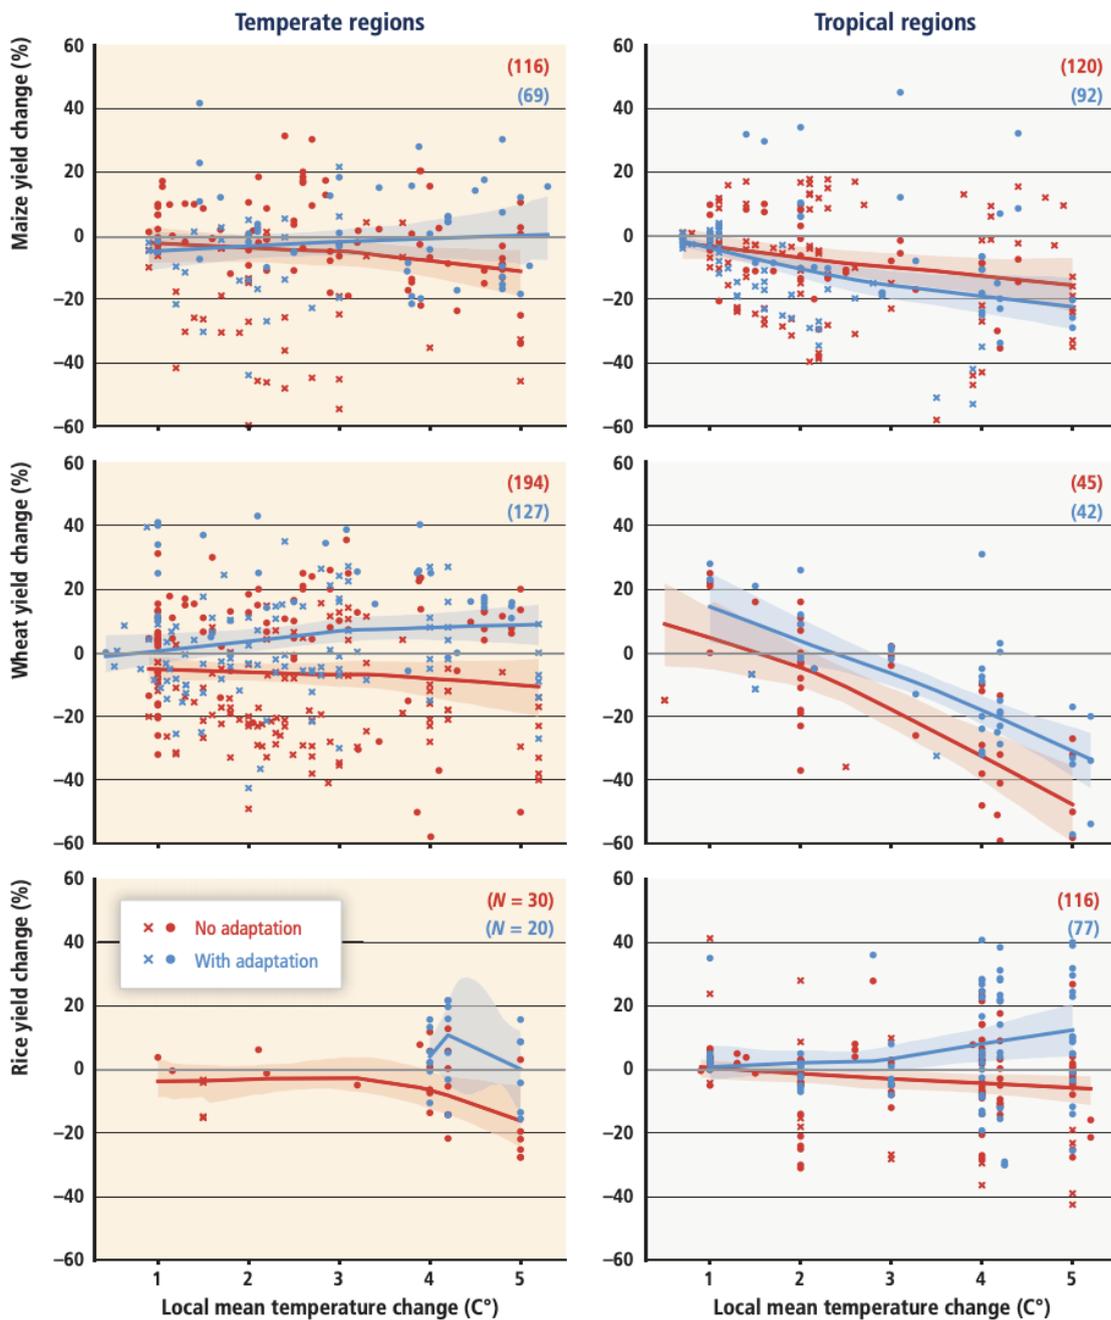
moisture and clouds in driving these trends (Christidis et al., 2005; Zwiers et al., 2011)." IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, Ch. 7, pp. 492-493.

¹⁹⁸ "Climate change is projected to reduce renewable surface water and groundwater resources significantly in most dry subtropical regions (robust evidence, high agreement). {3.4, 3.5} This will intensify competition for water among agriculture, ecosystems, settlements, industry, and energy production, affecting regional water, energy, and food security (limited evidence, medium to high agreement). {3.5.1, 3.5.2, Box CC-WE} In contrast, water resources are projected to increase at high latitudes" IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, Ch. 3 p. 232.

¹⁹⁹ "In high latitudes (such as Russia, northern Europe, Canada, South America) global warming may increase yields and expand the growing season and acreage of agricultural crops, although yields may be low due to poor soil fertility and water shortages in some regions." IPCC, *Climate Change 2014: Impacts, Adaptation, and Vulnerability*, Fifth Assessment Report (Cambridge University Press, 2014), Ch 7, p. 512.

²⁰⁰ "The model also inferred significant positive effects of precipitation (t =3.0; P =0.0031) and CO₂ (t =3.1; P =0.0022) with average yield increases of 0.53% (per % 1P), 0.06% (per ppm 1CO₂), respectively (Table 1)." A. J. Challinor et al., 'A Meta-Analysis of Crop Yield under Climate Change and Adaptation', *Nature Climate Change* 4, no. 4 (April 2014): 288, <https://doi.org/10.1038/nclimate2153>.

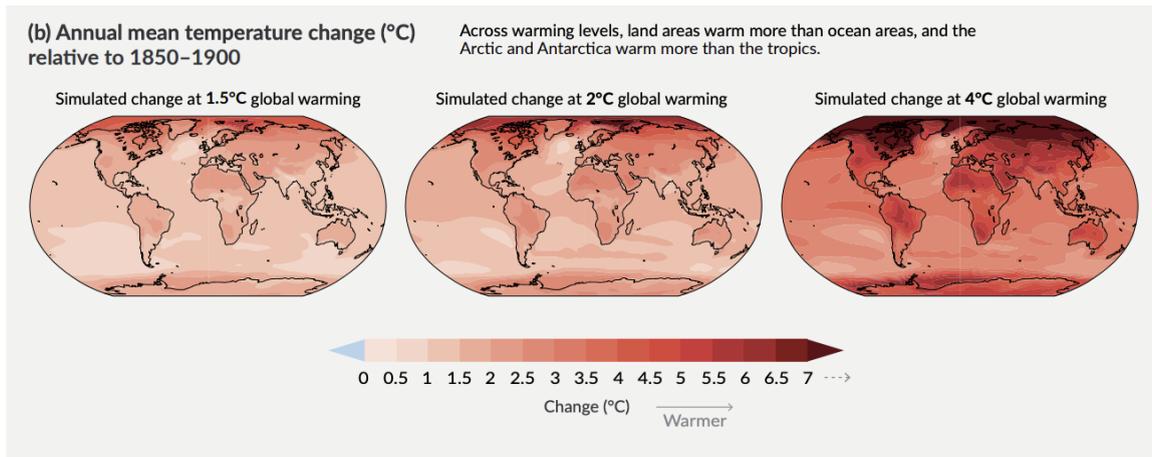
in the meta-analysis are “changes in varieties, planting times, irrigation and residue management”.²⁰¹



Source: IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, p. 498

When considering these effects, it is important to note that local warming can be different to global average warming, as shown in the figure below:

²⁰¹ A. J. Challinor et al., ‘A Meta-Analysis of Crop Yield under Climate Change and Adaptation’, *Nature Climate Change* 4, no. 4 (April 2014): 287, <https://doi.org/10.1038/nclimate2153>



Source: IPCC, *Climate Change 2021: The Physical Science Basis, Sixth Assessment Report* (UNFCCC, 2021), Summary for Policymakers, SPM.5.

At 4°C of warming, the Arctic is 7°C warmer. Warming is also greater on land than at sea, so land warming will typically be higher than average global surface warming, on average. For 4°C of average global warming, many highly populated regions will see local temperatures increase by 5–6°C.

Zhao et al (2017)

In a multi-model study of the effects of climate change on yields of the major crops, Zhao et al (2017) find that “Without CO₂ fertilization, effective adaptation, and genetic improvement, each degree-Celsius increase in global mean temperature would, on average, reduce global yields of wheat by 6.0%, rice by 3.2%, maize by 7.4%, and soybean by 3.1%.”²⁰² So, on these pessimistic assumptions, after 5°C of warming, yields would decrease by 15% to 35%. Once we relax these assumptions, the effects of climate change would be smaller.

Jägermeyr et al (2021)

Jägermeyr et al uses latest-generation crop and climate models to project future crop yields under different climate change scenarios. One major advantage the study has is that it does not just model the effects of temperature, but also precipitation changes, temperature-moisture feedbacks and CO₂ fertilisation.²⁰³ However, their models do not account for adaptation, or for other productivity improvements, which they leave to future work.²⁰⁴

²⁰² Chuang Zhao et al., ‘Temperature Increase Reduces Global Yields of Major Crops in Four Independent Estimates’, *Proceedings of the National Academy of Sciences* 114, no. 35 (29 August 2017): 9326–31, <https://doi.org/10.1073/pnas.1701762114>.

²⁰³ Recent literature has focused on capturing the temperature sensitivity of crops in isolation^{17–19}. To quantify climate change impacts more comprehensively, additional factors including precipitation changes, temperature–moisture feedbacks and [CO₂] need to be considered. The projections presented here dynamically respond to these climate drivers and shed new light on the effects of elevated [CO₂], which are among the largest sources of uncertainty in long-term crop yield estimates” Jägermeyr et al., ‘Climate Impacts on Global Agriculture Emerge Earlier in New Generation of Climate and Crop Models’, 2.

²⁰⁴ “Cropping system adaptation can substantially reduce and even outweigh adverse climate change impacts, for example, by switching to other crops⁵¹ or better-adapted varieties⁵². Integrated into ISIMIP’s wider cross-sector activities, GGCM will systematically evaluate farming system adaptation and changes in yield variability and extreme event impacts in subsequent efforts^{53,54}.” Jonas Jägermeyr et al., ‘Climate Impacts on Global Agriculture Emerge Earlier in New Generation of Climate

Their results are shown below:

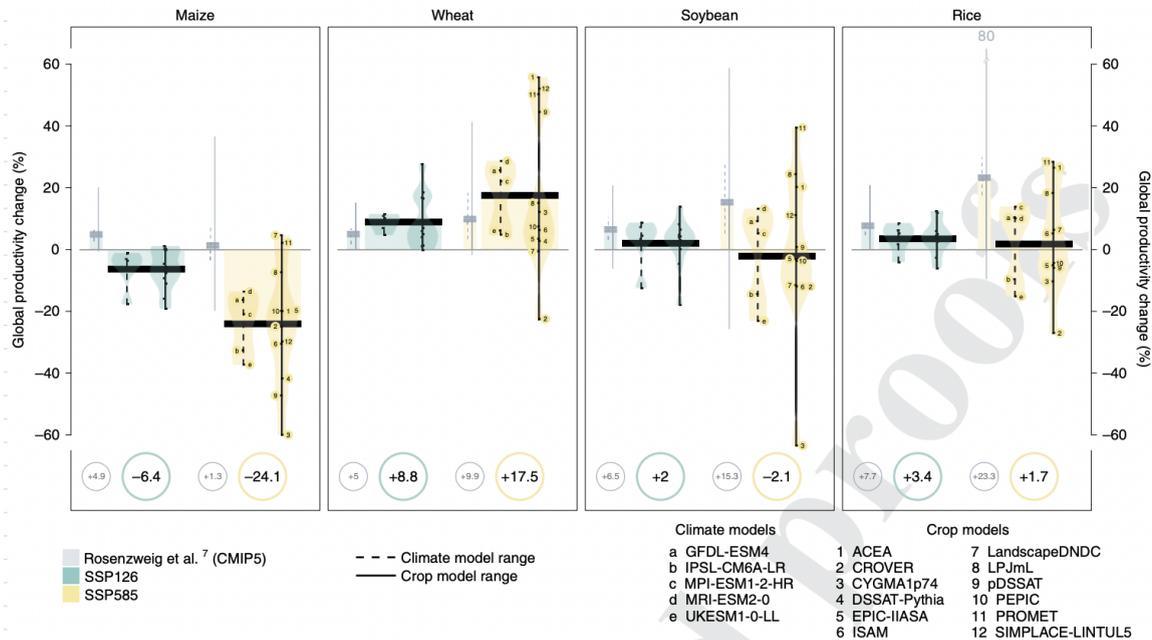


Fig. 1 | Ensemble end-of-century crop productivity response. Global productivity changes (2069–2099 compared to 1983–2013) for SSP126 and SSP585 are shown as the mean across climate and crop models for the four major crops (highlighted by bullets underneath the plot). Whiskers indicate the range of individual climate model realizations (dashed line, as the mean across crop models), and the range across crop models (solid line, as the mean across climate models). Individual model results are indicated by the bullets along the whisker lines (for SSP585 only); violin shades additionally highlight the model distribution. For context, grey bars and whiskers reference previous GGCM1 simulations based on CMIP5 (GC5; Rosenzweig et al.⁷) in the same way, without specifying individual models. Data are shown for the default [CO₂]. Not all crop models simulate all crops, see Supplementary Table 3 for details.

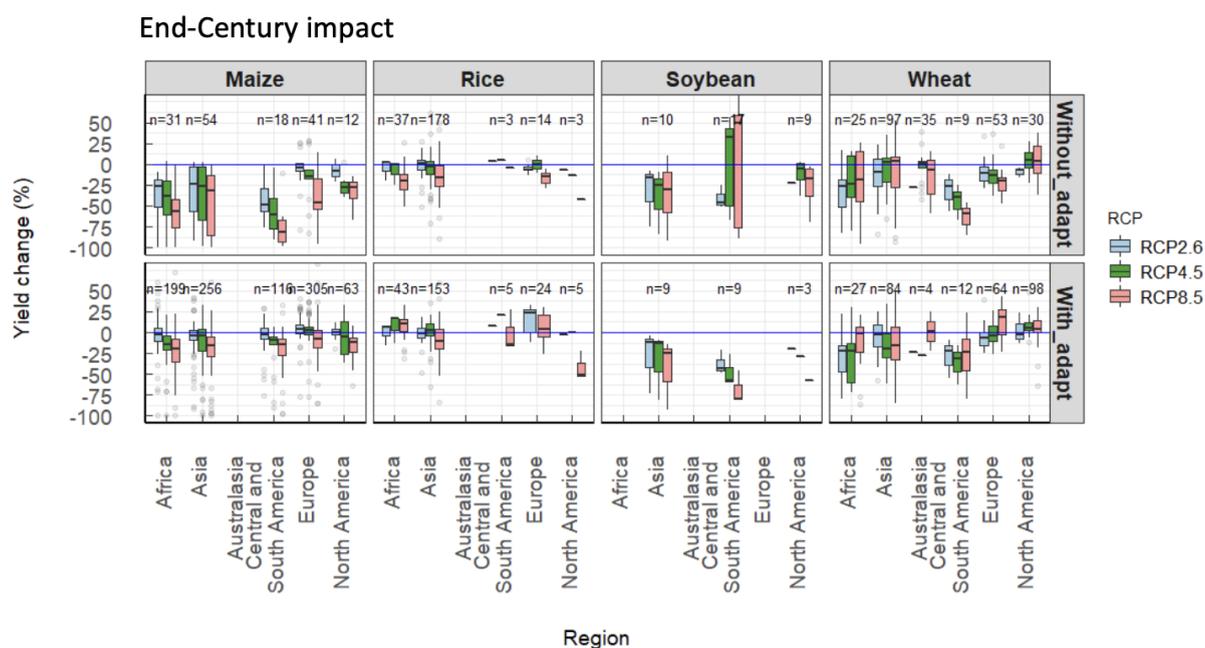
The overall picture is more pessimistic than found in earlier models around the time of the IPCC’s Fifth Assessment Report, though it is unclear how adaptation might affect their findings. On RCP8.5, on which there is 4.4°C warming by 2100, results for wheat are better than before, but soy, rice and maize results are worse.

IPCC Sixth Assessment Report

In the 2022 Sixth Assessment Report, the IPCC again notes that climate change is likely to have positive effects in temperate regions but negative effects in warmer regions.²⁰⁵ The IPCC mainly relies on studies by Hasegawa et al on crop yields. The table below shows the Hasegawa et al projection of the effect on climate change in different regions and on different emissions scenarios, Recall that RCP4.5 implies about 2.7°C, while RCP8.5 implies about 5°C of warming:

and Crop Models’, *Nature Food* 2, no. 11 (November 2021): 9, <https://doi.org/10.1038/s43016-021-00400-y>.

²⁰⁵ “The projected effects of climate change are positive where current annual mean temperatures (Tave) are below 10 °C, but they become negative with Tave above around 15°C. At Tave>20°C, even a small degree of warming could result in adverse effects.” IPCC, *Climate Change 2022: Impacts*, Ch. 5 sec. 5.4.3.2.



Supplementary Fig. S3 Climate change impacts on four crops in the mid 21st century with and without adaptation in IPCC regions by regions at mid-century (MC, 2040-2069, upper panels) and end-century (EC, 2070-2100, lower panels). n is the number of simulations. The box is the interquartile range (IQR) and the middle line in the box represents the median. The upper- and lower-end of whiskers are median $1.5 \times \text{IQR} \pm \text{median}$. Open circles are values outside the $1.5 \times \text{IQR}$.

Source: Toshihiro Hasegawa et al., 'A Global Dataset for the Projected Impacts of Climate Change on Four Major Crops', *Scientific Data* 9, no. 1 (16 February 2022)
<https://doi.org/10.1038/s41597-022-01150-7>.

This shows that there is scope for adaptation to significantly dampen the effects of climate change, especially on maize yields.

By adaptation Hasegawa et al mean any changes from current methods in terms of "fertiliser, irrigation, cultivar, soil organic matter management, planting time, tillage, and others".²⁰⁶

Correlated yield declines

As well as projections of overall likely trends in average yields, it is useful to assess the risk of synchronised declines in yields in multiple crops, due to increasing yield variability and to extreme weather events. Crop losses in a single, main crop producing region can be offset through trade with other crop-producing regions. If several breadbaskets suffer from negative climate impacts at the same time, however, the effects could be more substantial.

Tigchelaar et al. (2018)

²⁰⁶ Hasegawa et al., 'A Global Dataset for the Projected Impacts of Climate Change on Four Major Crops'.

Tigchelaar et al. (2018) explore the risk that the currently top four maize producing countries experience synchronised production declines of more than 10% relative to the average today, warming scenarios. Their findings are as follows, assuming that there are no improvements in the heat tolerance of crops.²⁰⁷

- Today, the risk of such a synchronised decline in any given year is virtually zero
- For 4°C, the risk increases to 86% in any given year.

The following chart shows the findings of the effect of Tigchelaar et al (2018) on the probability of different yield outcomes in different countries. The red line shows the probability distribution across yields for 4°C of warming, while the black line shows the probability distribution across yields for the present day.

²⁰⁷ “We find that as the global mean temperature increases, absent changes in temperature variability or breeding gains in heat tolerance, the coefficient of variation (CV) of maize yields increases almost everywhere to values much larger than present-day values.” Michelle Tigchelaar et al., ‘Future Warming Increases Probability of Globally Synchronized Maize Production Shocks’, *Proceedings of the National Academy of Sciences* 115, no. 26 (26 June 2018): 6644–49, <https://doi.org/10.1073/pnas.1718031115>.

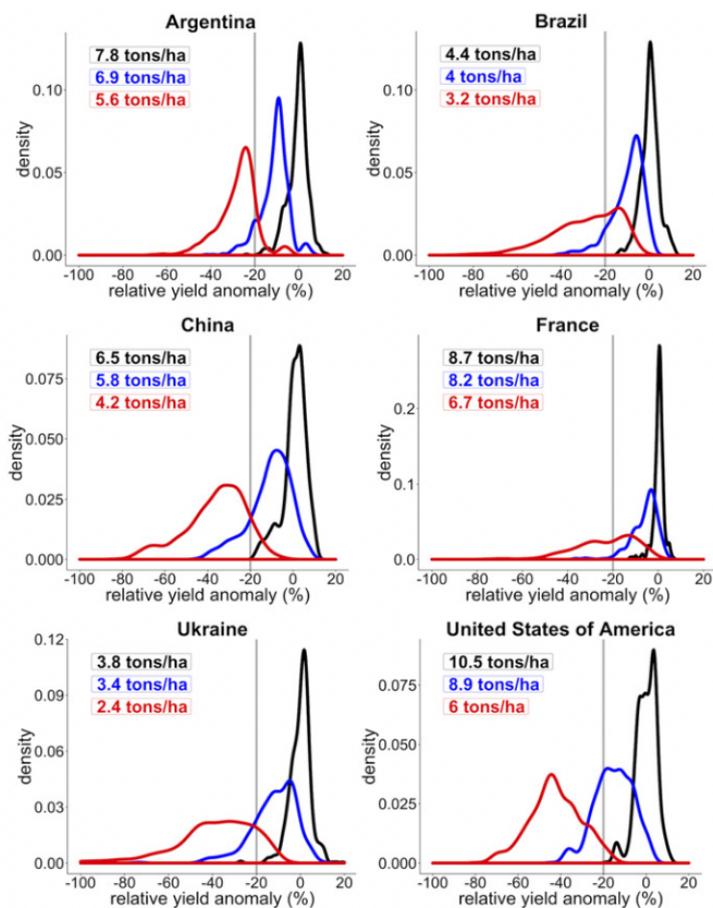


Fig. 4. Warming-induced changes in yield variability in top-producing regions of the six largest maize-producing and -exporting countries (*SI Appendix, Table S2*): Probability density functions of yield anomalies with respect to present-day mean yield for present-day climate (black), following 2 °C of annual mean global warming (blue), and following 4 °C of annual mean global warming (red). The vertical gray line denotes a relative yield reduction of 20%, and boxed values indicate mean present-day yield in these areas for present-day climate (1999–2008; black) and for 2 °C (blue) and 4 °C (red) warming.

As this shows, Tigchelaar et al (2018) assume that for 4°C of warming, average yield will be lower than today. However, this assumption is not realistic because there will be countervailing improvements in agricultural productivity in the time it takes to reach 4°C. I discuss this point in more detail in section 5.4. Over the last 60 years, crop yields have increased by upwards of 200-300%, and it is reasonable to think that they will increase by a further 100% in the next 80 years. So, overall yields would still be much higher than today, even if there were synchronised declines by 20% in major food producing regions.

Gaupp et al (2019)

Gaupp et al (2019) assesses the risk of synchronised large declines in yield for multiple crops. Specifically, they assess the probabilities of events when the climatic conditions are at

least as bad as occurred when historical yields were in the 25th percentile of yields in the study period.²⁰⁸

They found that the return rate of synchronised low yields in all five breadbaskets increases at 2°C.²⁰⁹

- Wheat: Historical return rate is 43 years vs. 15 years under 2°C.
- Maize: Historical return rate is 16 years vs. <2 years under 2°C.
- Soybean: Historical return rate is 20 years vs. 9 years under 2°C.

Production losses increase at 2°C relative to 1.5°C. If multiple breadbasket failures occur, the losses would be:

- Wheat: 8.6 million tonnes (or 1% of 2018 production)²¹⁰
- Maize: 19.8 million tonnes (or 2% of 2018 production)²¹¹
- Soybean: 9.9 million tonnes (or 3% of 2018 production)²¹²

For comparison:

“Historical examples of global crop production shocks include 7.2 million tons soybean losses in 1988/99 and 55.9 million tons maize losses in 1988 which were mostly caused by low rainfall and high temperatures during summer growing season in the US.”²¹³

Adjusting for the probabilities of these losses under 2°C, *expected* the losses are:

- Wheat: 161,000 tonnes (or 0.02% of 2018 production)
- Maize: 2,753,000 tonnes (or 0.3% of 2018 production)
- Soybean: 265,000 tonnes (or 0.08% of 2018 production)

²⁰⁸ “We identify a ‘breadbasket failure’ event as being when the climatic conditions are at least as severe as those conditions associated with the 25 percentile of the logistically detrended yields (with detrended yields as residuals of the non-linear logistic regression with a residual mean equal to zero). The crop production loss for an event of this severity is the 25 percentile of the logistically detrended yield multiplied with the 2012 harvested area” Franziska Gaupp et al., ‘Increasing Risks of Multiple Breadbasket Failure under 1.5 and 2 C Global Warming’, *Agricultural Systems* 175 (2019): 40.

²⁰⁹ “For wheat, which shows the smallest simultaneous climate risks, the return period for all five breadbaskets exceeding their climate thresholds decreases from 43 years (or 0.023 annual probability under historical conditions to 21 years (0.047) in a 1.5 °C scenario and further down to around 15 years (0.066) under 2 °C. Soybean has a return period of simultaneous climate risks in all breadbaskets of around 20 years (0.049 today which decreases to 9 (0.116) and 7 years (0.143) in a 1.5 and 2 °C warmer world respectively. Maize risks are highest in our study with an initial return period of 16 years (0.061), decreasing to < 3 (0.39) and < 2 years (0.538) under future global warming.” Gaupp et al., ‘Increasing Risks of Multiple Breadbasket Failure under 1.5 and 2 C Global Warming’.

²¹⁰ Wheat production in 2018 was 735 million tonnes ([Our World in Data](#)).

²¹¹ Maize production in 2018 was 1.1 billion tonnes ([Our World in Data](#)).

²¹² Soy production in 2018 was 348 million tonnes ([Our World in Data](#)).

²¹³ Gaupp et al., ‘Increasing Risks of Multiple Breadbasket Failure under 1.5 and 2 C Global Warming’, 42.

As I discuss below, these expected effects are small relative to total production and relative to likely improvements in agricultural production.

5.3.2. Crop migration

The literature discussed above assumes that crop migration is not available as an adaptive response, which likely leads it to overestimate the effect on global agricultural production.

“Many studies have estimated the adverse effects of climate change on crop yields, however, this literature almost universally assumes a constant geographic distribution of crops in the future. Movement of growing areas to limit exposure to adverse climate conditions has been discussed as a theoretical adaptive response but has not previously been quantified or demonstrated at a global scale”²¹⁴

Sloat et al (2020) argue that crop migration has mitigated the effects of warming on yield since 1973. Blomqvist et al (2020) argue that crop migration accounted for around one sixth of the global increase in crop production since 1960.²¹⁵

According to Zabel et al (2014), with no adaptation and on A1B (a high emissions scenario similar to RCP8.5), the land area suitable for agriculture is higher by 2100 mainly due to gains in China, the US, Russia and Canada, though the tropics and subtropics lose out.

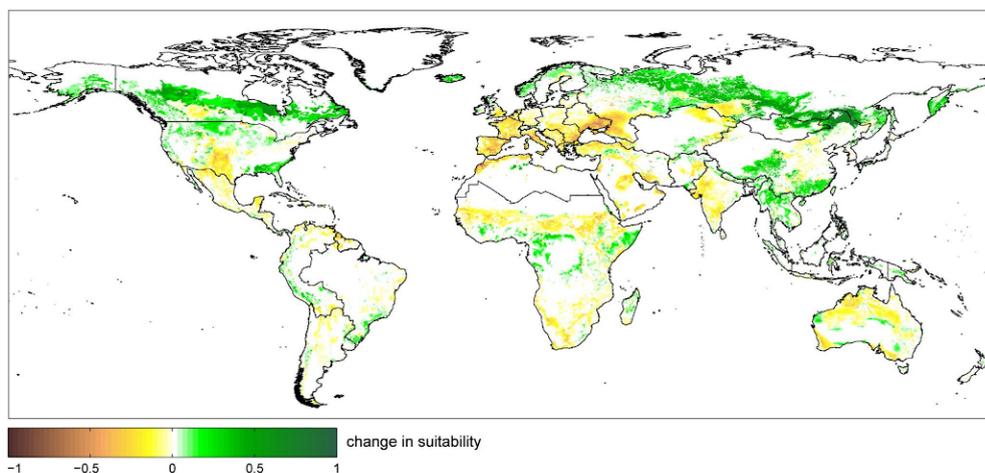


Figure 11. Change in agricultural suitability between 1981–2010 and 2071–2100. Green areas indicate an increase in suitability while brown areas show a decreasing suitability.
doi:10.1371/journal.pone.0107522.g011

Source: Florian Zabel, Birgitta Putzenlechner, and Wolfram Mauser, ‘Global Agricultural Land Resources – A High Resolution Suitability Evaluation and Its Perspectives until 2100 under Climate Change Conditions’, *PLOS ONE* 9, no. 9 (17 September 2014): e107522, <https://doi.org/10.1371/journal.pone.0107522>.

²¹⁴ Lindsey L. Sloat et al., ‘Climate Adaptation by Crop Migration’, *Nature Communications* 11, no. 1 (6 March 2020): 1243, <https://doi.org/10.1038/s41467-020-15076-4>.

²¹⁵ ‘Country share’ describes the geographic distribution of cropland. A shift in cropland from lower-yielding to higher-yielding countries, for instance, would boost aggregate yield without any one country improving its yields. Linus Blomqvist, Luke Yates, and Barry W. Brook, ‘Drivers of Increasing Global Crop Production: A Decomposition Analysis’, *Environmental Research Letters* 15, no. 9 (September 2020): fig. 2, <https://doi.org/10.1088/1748-9326/ab9e9c>.

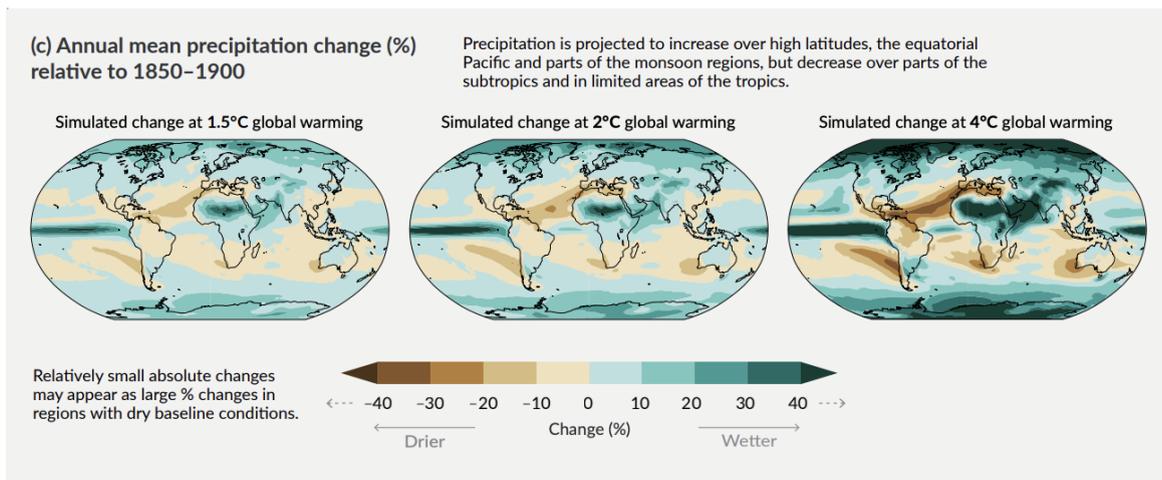
This suggests that using the meta-analyses of the effect of climate change on yield that only account for in-situ adaptation will tend to overestimate the damage of climate change to global agriculture.

5.3.3. Water stress and agriculture

The effect of climate change on precipitation

My discussion here relies heavily on the overview of climate change on precipitation by Carbonbrief [here](#).

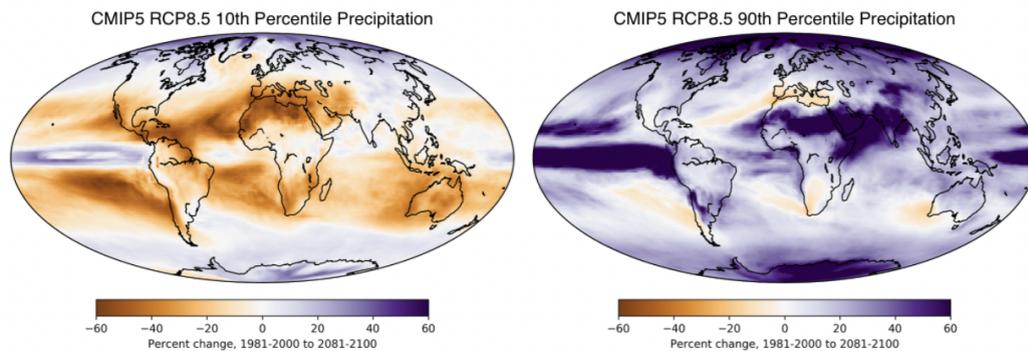
The figure below shows projected percentage change in precipitation at different levels of warming:



Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Summary for Policymakers, SPM.5.

Substantial parts of Latin America, the southern USA, southern Africa, the Mediterranean, and Australia are projected to dry out in the future. The Mediterranean region is expected to have around 20% less precipitation by 2100 in an RCP8.5 world, with similar reductions also found in southern Africa. Western Australia, Chile, and Central America/Mexico may all become around 10% drier. These changes tend to increase with warming.

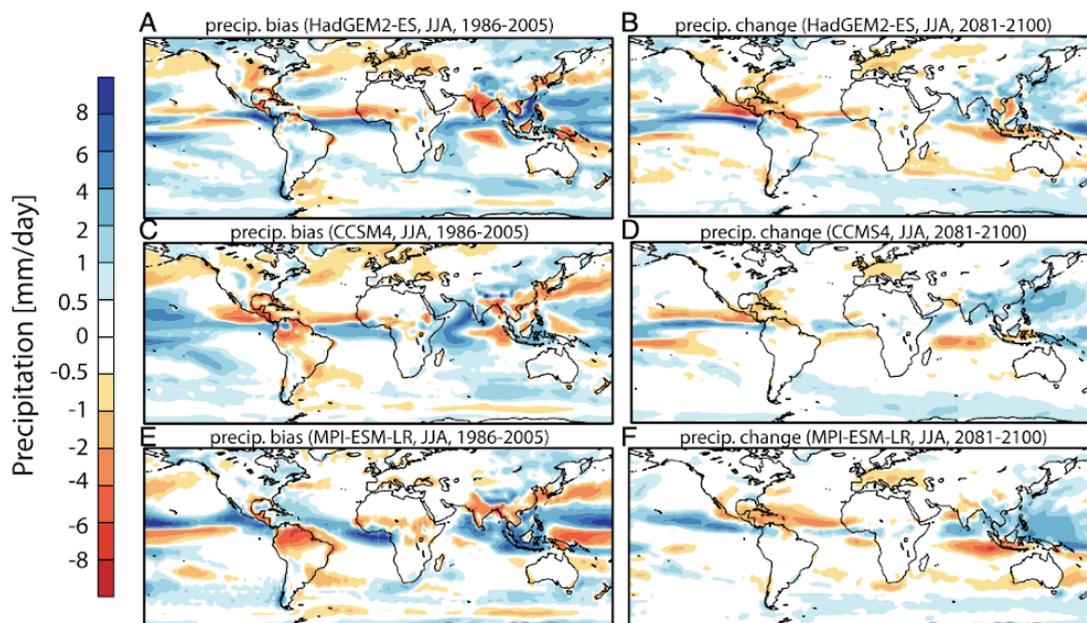
However, there is considerable disagreement among models about the effect of climate change on precipitation. The figure below shows the driest projection and wettest projections for each different part of the world across all the CMIP5 models, represented by the 10th and 90th percentile of all the models (i.e. the 10% of models that show the most reduction in precipitation and the 10% that show the most increase in precipitation for any region of the world).



RCP8.5 10th percentile of mean precipitation change (left map) and 90th percentile (right map) for total precipitation (rain and snow) for each 1×1 latitude/longitude gridcell between 1981-2000 and 2081-2100. Uses one run for each model, 38 models total. Data from [KNMI Climate Explorer](#); maps by Carbon Brief.

Source: [Carbonbrief](#)

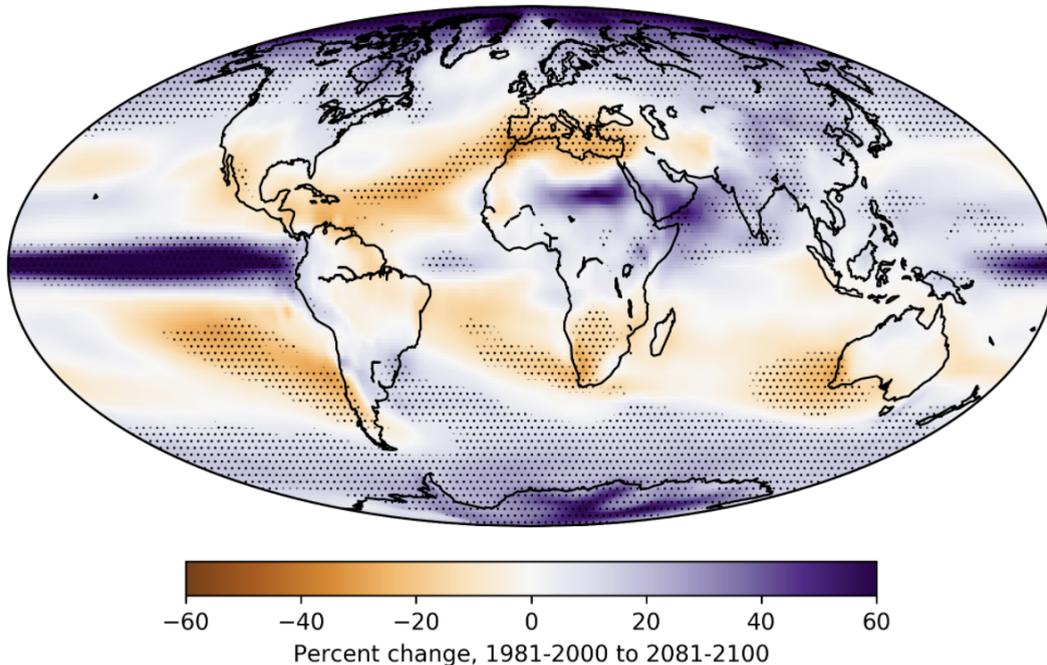
Climate models are poor at reproducing observed changes in precipitation. The left-hand panes of the figure below show the systematic error of three of the leading models used in the last IPCC report for precipitation in June, July and August for 1986 to 2005. The right-hand pane shows the projected effect of RCP8.5 on precipitation by 2100 on the same models. There are widespread regions where the systematic error exceeds the climate change signal (in some regions by more than a factor of 20)



Source: Tim Palmer and Bjorn Stevens, 'The Scientific Challenge of Understanding and Estimating Climate Change', *Proceedings of the National Academy of Sciences* 116, no. 49 (3 December 2019): 24390–95, <https://doi.org/10.1073/pnas.1906691116>.

There is some agreement across models on projected effects on precipitation. The figure below shows the same annual average change in precipitation between today and the end of the century, but adds dots to indicate areas where at least nine out of 10 models agree on the direction of change.

CMIP5 RCP8.5 multimodel mean all precipitation



As first figure, but with areas where 90% of the models agree on the sign of the change highlighted with dots. Data from [KNMI Climate Explorer](#); map by Carbon Brief.

Source: [Carbonbrief](#)

There is widespread agreement among the models that both the tropical Pacific and high-latitude areas will have more precipitation in the future. India, Bangladesh and Myanmar will all become wetter, as will much of northern China.

The models largely agree that the Mediterranean region and southern Africa will have less precipitation in the future. They also agree on reduced precipitation in southwest Australia around Perth, in southern Chile, the west coast of Mexico and over much of the tropical and subtropical Atlantic ocean.

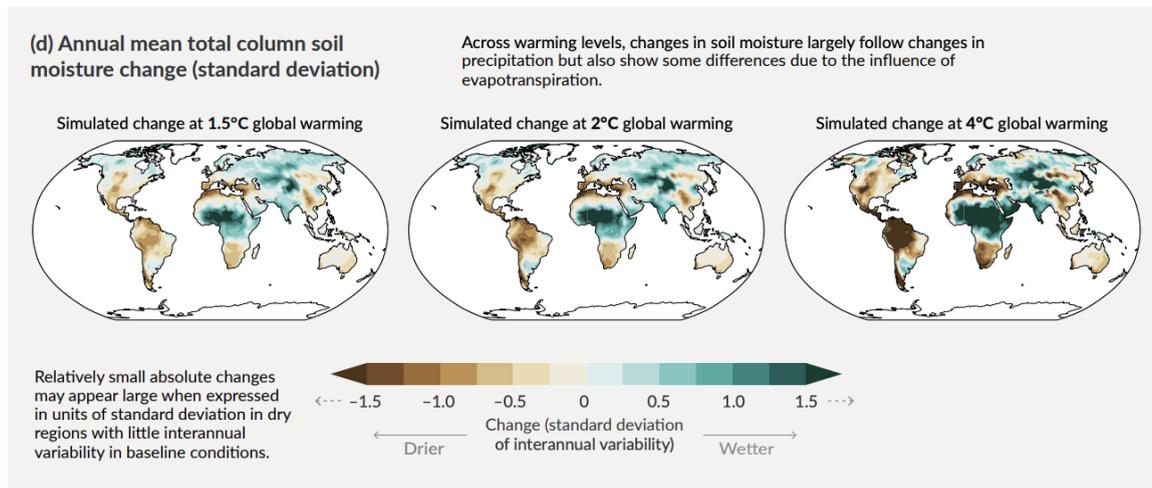
Given the problems climate models have at reproducing observed precipitation, it is difficult to put much weight on regional projections even where there is agreement among 90% of models. For example, the models agree that precipitation will increase across much of India on RCP8.5, but then we saw above that three leading climate models made huge errors in reproducing *observed* precipitation in India, *and* the sign of the error was different across the three models. Therefore, even if 90% of the models agree on the sign of a change, there seems to be a reasonable chance that they are all wrong.

As I mentioned in Chapter 2, this uncertainty means we should be *more* worried about climate change, not less.

The effect of climate change on drought

The effect of climate change on drought is determined not only by changes in precipitation, but also by higher temperatures, which increase the rate of evaporation and soil moisture loss.

The maps below show the effects on soil moisture at different levels of warming, according to climate models.



Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Summary for Policymakers, SPM.5.

What effect will climate-induced water stress have on agriculture and the economy?

Changes in precipitation and drought create challenges for agriculture. Increased flooding might affect crop production, but the greatest risk comes from increased aridity. As we have seen, it is difficult to have much confidence in projections of which regions will dry out. Still we know that climate change will cause some areas to dry out.

Irrigation is an effective adaptation measure

In the meta-analysis used in the IPCC's Fifth Assessment report, shown above, increased irrigation was found to be an effective adaptation to declining water availability.²¹⁶ The report cites three studies which project zero or modest (~15%) global increases in demand for irrigation by 2100, with some regions, such as Europe, the USA and parts of Asia, experiencing large increases (>40%), on some models.²¹⁷

²¹⁶ "Supplementary Table 4 200 shows that of the four categories (for this analysis fertiliser was included with "other") only irrigation and "other", on average, increase yields from baselines values. Of the four categories, irrigation is the one that is most likely to systematically increase yields, since planting date and cultivar changes can reduce yields." A. J. Challinor et al., 'A Meta-Analysis of Crop Yield under Climate Change and Adaptation', *Nature Climate Change* 4, no. 4 (April 2014): SI p. 5, <https://doi.org/10.1038/nclimate2153>.

²¹⁷ "Using projections from 19 CMIP3 GCMs forced by SRES A2 emissions to drive a global vegetation and hydrology model, climate change by the 2080s would hardly alter the global irrigation water demand of major crops in areas currently equipped for irrigation (Konzmann et al., 2013). However, there is high confidence that irrigation demand will increase significantly in many areas (by more than 40% across Europe, USA, and parts of Asia). Other regions—including major irrigated areas in India, Pakistan, and southeastern China—might experience a slight decrease in irrigation demand, due for example to higher precipitation, but only under some climate change scenarios (also see Biemans et al., 2013). Using seven global hydrological models but a limited set of CMIP5

The economic barriers to increasing irrigation to adapt to climate change seem low. 35% of agricultural land in India is irrigated, but income per head at purchasing power parity is only \$6,400, compared to the global average of \$17,000. In Bangladesh, 60% of agricultural land is irrigated. Farmers in the Fertile Crescent at the dawn of agriculture made extensive use of irrigation.²¹⁸ In the future, when most countries will be richer and have access to better technology, they are also likely to make use of irrigation if it is beneficial.

For low growth agrarian economies, the effects on water stress could have bad humanitarian consequences. The people who have contributed the least to climate change would be hardest hit. We have strong reasons to reduce emissions for the sake of these people.

Water management is poor

Water is very poorly managed throughout the world. Cheap or free water is one of innumerable implicit and explicit subsidies that farmers receive at the expense of society and the environment. Damania (2020) notes that

“Most water is allocated to agriculture and much is made available at no cost to the user. As a result, overuse of water is a common problem, especially in the agricultural sector. The marginal value of water in different uses varies a great deal because the prices paid by industry, agriculture, and residential users often have no relation to the opportunity costs of supplying water to them. As an example, Olmstead (2010) notes that in the desert state of Arizona in the United States, water prices vary from \$27/acre-foot for agriculture to \$3,200/acre-foot for urban uses. While some of the variation can be explained by the difference in the quality of the product being delivered, most of it is a function of market and institutional failures that do not allocate water based on its economic value. This suggests that the benefits of reallocating water from farms to cities would be large.”

The potential benefits of market pricing of water are large, on the order of a 6% boost to global GDP, on some estimates.²¹⁹

projections, Wada et al. (2013) suggested a global increase in irrigation demand by the 2080s (ensemble average 7 to 21% depending on emissions scenario), with a pronounced regional pattern, a large inter-model spread, and possible seasonal shifts in crop water demand and consumption. By contrast, based on projections from two GCMs and two emissions scenarios, a slight global decrease in crop water deficits was suggested in both irrigated and rainfed areas by the 2080s, which can be explained partly by a smaller difference between daily maximum and minimum temperatures (Zhang and Cai, 2013). As in other studies, region-to-region variations were very heterogeneous.” IPCC, *Climate Change 2022: Impacts*, Ch. 7.

²¹⁸ “Throughout the region, irrigation is necessary for the best agricultural results and, indeed, is often essential to any farming at all. Radiocarbon dating has shown that incipient agriculture and village agglomerations in the Fertile Crescent there must be dated back to about 8000 BCE, if not earlier, and that the use of irrigation followed rapidly.” [Britannica](#)

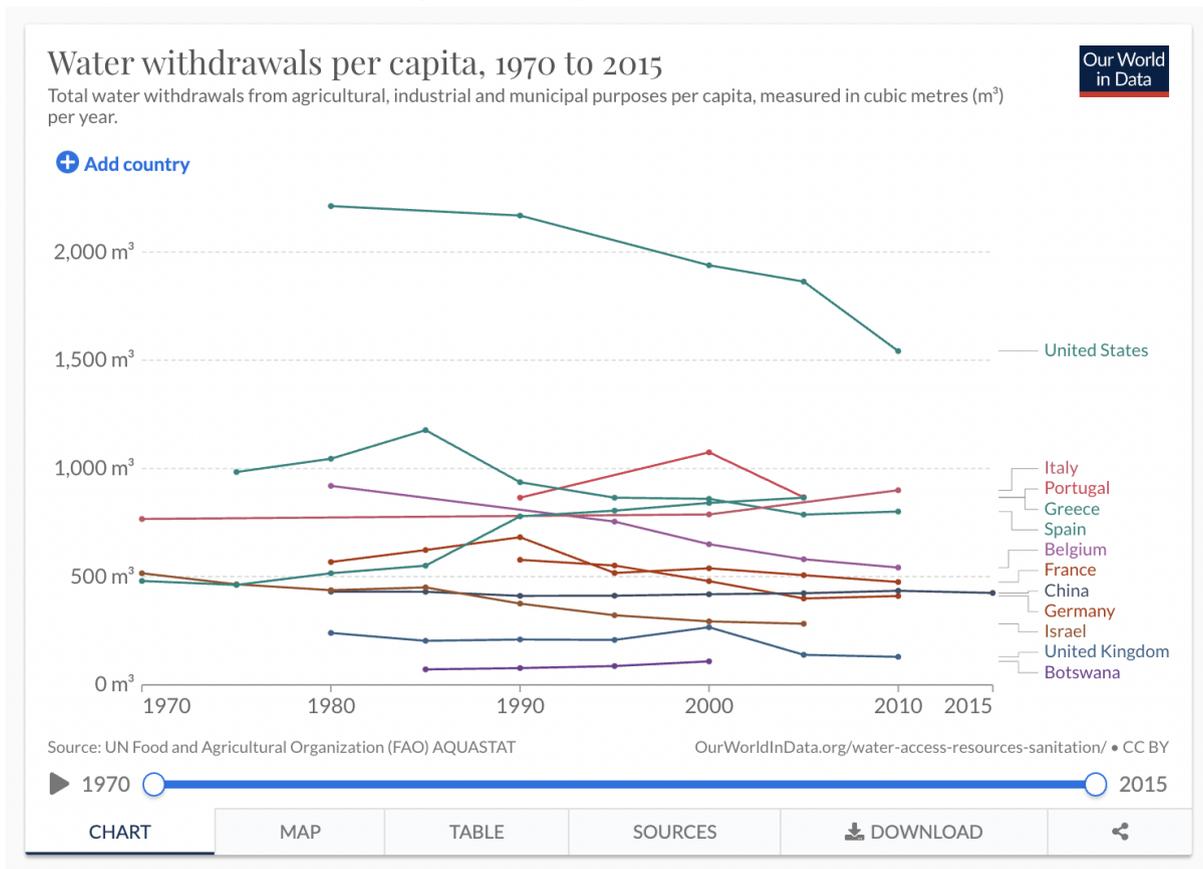
²¹⁹ “In a global context Roson and Damania (2017) find losses of about 6 per cent of GDP. Their simulations suggest that even if only a part of water use is allocated based on its economic price that brings supply and demand into balance, much of the distortionary impact vanishes” Richard Damania, ‘The Economics of Water Scarcity and Variability’, *Oxford Review of Economic Policy* 36, no. 1 (6 January 2020): 24–44, <https://doi.org/10.1093/oxrep/grz027>.

Desalination

Another more expensive adaptive response to drought and declining freshwater availability is desalination. There is a good review of trends in desalination use and cost in [this review](#) by the Advisian Worley Consulting Group.

From 1960 to 2010, desalination costs fell by a factor of 10 from \$10/m³ of water to less than \$1/m³. By 2017, costs had declined by a further 20%.

The costs to fill any shortfalls in water consumption through desalination seem manageable for wealthy countries. Consider this illustrative example. Total Spanish per capita water withdrawals per year, including agricultural, industrial and municipal uses, are around 1,000m³. This is above the European average.



To supply the entirety of Spanish per capita water consumption through desalination would cost less than \$1,000, compared to current GDP per capita of \$30,000. Moreover, it is unrealistic to imagine that all Spanish water consumption would need to be supplied by desalination - much of the supply would still come from freshwater sources.

If the trend in cost declines of desalination continues, then desalination costs would fall by a factor of 10 by 2070. Thus, the annual costs to supply each Spaniard's annual water use would fall to \$100. However, two experts have told me that we are already close to fundamental physical limits on desalination.

Some water-stressed countries already deploy desalination at a large scale. For example, Israel started [a desalination programme in 2000](#) to deal with water stress. It now desalinates 750 million m³ of water per year, which is 83m³ per person, or around 30% of total per person water use.

There would be strong incentives to use desalination if the supply of freshwater ever becomes a problem. As before, poor agrarian economies would likely be unable to afford desalination, so for them adaptation would be much harder.

What are the economic costs of water stress?

In recent years, innovative studies using panel data have been used to estimate the effect of weather on GDP. The basic idea is that an average year in a particular country is the ‘control’ and a dry year is ‘the treatment group’. The same studies have been used to estimate the future effects of climate change.

One problem that these studies have is that interannual weather variation is different to long-term climate change, a point I expand in Chapter 10. It is not clear in which direction this biases the estimates.

The highest profile papers by Burke et al (2015a) and Dell et al (2012) conclude that there is no robust evidence that rainfall variation has a consistent and negative effect on GDP growth, though they find that temperature variation does have detrimental effects. I discuss these papers in more detail in Chapter 10. Other papers produce mixed results.

Table 1: Panel regressions of changes in water availability (rainfall) on measures of aggregate economic activity

	GDP growth	GDP level	Agricultural GDP growth
Burke <i>et al.</i> , 2015	No significant impact		
Dell <i>et al.</i> , 2012	No significant impact		No significant impact
Brown <i>et al.</i> , 2013	Positive significant impact		No significant impact
Sadoff <i>et al.</i> , 2015			Positive significant impact
Barbier, 2015	Negative significant impact		
El Khanji and Hudson, 2016	No significant impact	Significant impact	

Note: A blank cell implies that regressions on that dependent variable have not been reported in the paper.

Source: Richard Damania, ‘The Economics of Water Scarcity and Variability’, *Oxford Review of Economic Policy* 36, no. 1 (6 January 2020): 24–44, <https://doi.org/10.1093/oxrep/grz027>.

A study of the effect of future drought on Europe by Naumann et al (2021) found that, with adaptation, drought conditions at 4°C of warming above pre-industrial levels would reduce GDP in the EU and UK by 0.07% per year.

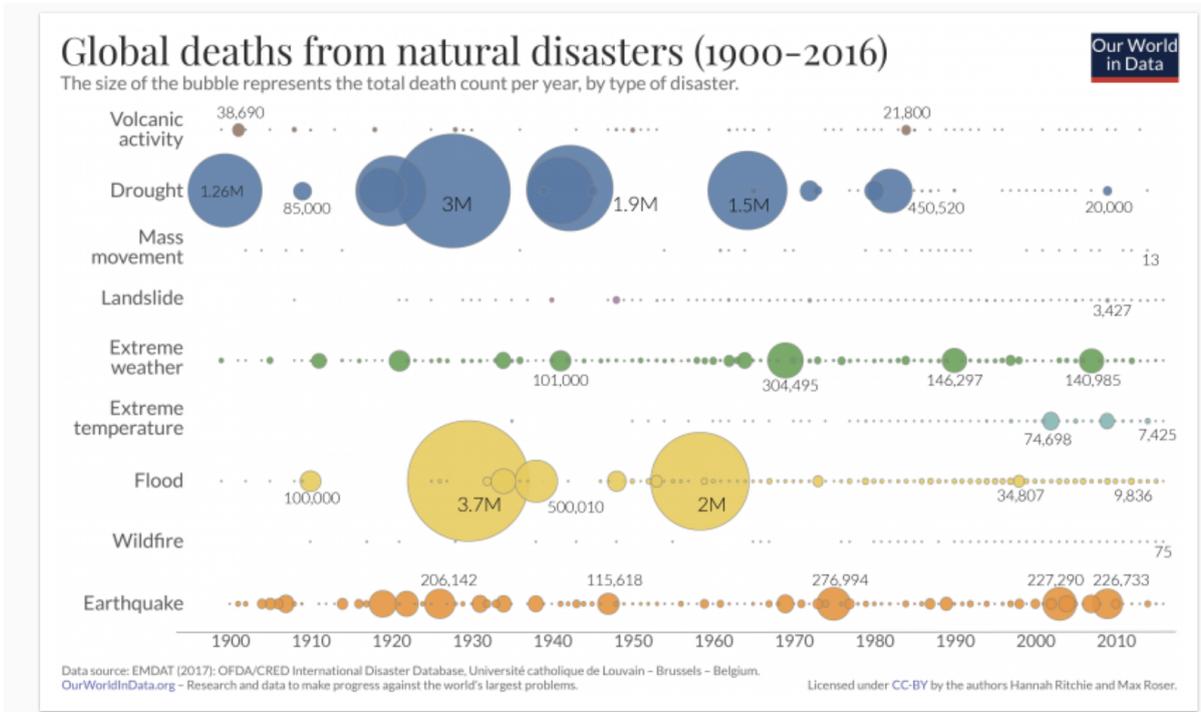
Country	Base economy					Economy 2100 static vulnerability				Economy 2100 dynamic vulnerability			
	base	1.5°C	2.0°C	3.0°C	4.0°C	1.5°C	2.0°C	3.0°C	4.0°C	1.5°C	2.0°C	3.0°C	4.0°C
Austria	0.08	0.06	0.05	0.03	0.05	0.04	0.04	0.02	0.04	0.02	0.02	0.01	0.02
Belgium	0.06	0.07	0.11	0.18	0.33	0.06	0.09	0.14	0.27	0.03	0.05	0.08	0.14
Bulgaria	0.24	0.32	0.36	0.47	0.88	0.23	0.26	0.35	0.65	0.10	0.11	0.15	0.28
Croatia	0.16	0.19	0.17	0.15	0.30	0.14	0.13	0.11	0.22	0.06	0.06	0.05	0.10
Cyprus	0.23	0.34	0.48	1.00	1.32	0.20	0.27	0.58	0.76	0.09	0.13	0.27	0.35
Czechia	0.11	0.08	0.07	0.07	0.07	0.05	0.05	0.05	0.05	0.02	0.02	0.02	0.02
Denmark	0.05	0.06	0.04	0.06	0.03	0.04	0.03	0.04	0.02	0.02	0.01	0.02	0.01
Estonia	0.09	0.05	0.03	0.03	0.01	0.03	0.02	0.02	0.01	0.01	0.01	0.01	0.00
Finland	0.04	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.01	0.00
France	0.06	0.05	0.11	0.18	0.31	0.04	0.07	0.13	0.21	0.02	0.04	0.06	0.11
Germany	0.04	0.03	0.03	0.04	0.04	0.02	0.03	0.03	0.03	0.01	0.01	0.02	0.02
Greece	0.16	0.25	0.32	0.44	0.92	0.20	0.24	0.34	0.70	0.09	0.12	0.16	0.34
Hungary	0.16	0.14	0.13	0.12	0.17	0.11	0.10	0.10	0.13	0.05	0.05	0.05	0.06
Ireland	0.07	0.12	0.17	0.34	0.30	0.10	0.15	0.30	0.26	0.05	0.07	0.15	0.13
Italy	0.09	0.11	0.11	0.14	0.23	0.08	0.09	0.11	0.18	0.04	0.05	0.06	0.09
Latvia	0.09	0.05	0.02	0.02	0.01	0.03	0.01	0.01	0.01	0.01	0.01	0.01	0.00
Lithuania	0.10	0.04	0.01	0.00	0.00	0.03	0.01	0.00	0.00	0.01	0.00	0.00	0.00
Luxembourg	0.05	0.01	0.05	0.06	0.12	0.01	0.03	0.04	0.08	0.00	0.02	0.02	0.05
Malta	0.19	0.22	0.23	0.29	0.48	0.15	0.16	0.20	0.33	0.07	0.08	0.09	0.16
Netherlands	0.04	0.05	0.09	0.12	0.16	0.04	0.07	0.10	0.13	0.02	0.03	0.05	0.06
Poland	0.09	0.05	0.04	0.02	0.01	0.03	0.03	0.02	0.01	0.02	0.01	0.01	0.00
Portugal	0.15	0.25	0.30	0.33	0.53	0.20	0.24	0.27	0.43	0.10	0.12	0.13	0.21
Romania	0.30	0.32	0.34	0.48	0.57	0.23	0.25	0.34	0.41	0.10	0.11	0.15	0.18
Slovakia	0.12	0.12	0.11	0.10	0.08	0.09	0.08	0.07	0.06	0.04	0.03	0.03	0.03
Slovenia	0.13	0.26	0.15	0.17	0.39	0.19	0.11	0.12	0.29	0.09	0.05	0.06	0.14
Spain	0.14	0.21	0.27	0.38	0.51	0.15	0.19	0.27	0.36	0.07	0.09	0.13	0.17
Sweden	0.04	0.01	0.00	0.01	0.01	0.01	0.00	0.00	0.01	0.00	0.00	0.00	0.00
United Kingdom	0.04	0.04	0.05	0.07	0.09	0.03	0.03	0.06	0.06	0.01	0.02	0.03	0.03
Mediterranean	0.12	0.16	0.19	0.26	0.40	0.12	0.14	0.19	0.29	0.06	0.07	0.09	0.14
Atlantic	0.05	0.05	0.08	0.14	0.21	0.04	0.06	0.10	0.15	0.02	0.03	0.05	0.08
Continental	0.06	0.05	0.05	0.06	0.07	0.04	0.04	0.05	0.05	0.02	0.02	0.02	0.03
Boreal	0.05	0.02	0.01	0.01	0.01	0.01	0.00	0.01	0.01	0.00	0.00	0.00	0.00
EU+UK	0.07	0.08	0.10	0.14	0.20	0.06	0.07	0.10	0.15	0.03	0.03	0.05	0.07

Table S7. Expected Annual Damage (EAD) as % of GDP for the baseline and different scenarios.

Source: Gustavo Naumann et al., 'Increased Economic Drought Impacts in Europe with Anthropogenic Warming', *Nature Climate Change* 11, no. 6 (June 2021): 485–91, <https://doi.org/10.1038/s41558-021-01044-3>.

We saw above that the Mediterranean would be one of the regions affected the worst by increasing droughts. According to Naumann et al (2021), there, the damage would rise to 0.14% of GDP per year by 2100, with adaptation. Low growth agrarian economies in the tropics and subtropics would be hit harder due to their limited ability to adapt. This suggests that the effects of climate change on water availability would be bad, but still very far from global agricultural collapse.

Given all these lines of evidence, water stress seems to be a weak lever on the risk of a global agricultural catastrophe. Indeed, over the 20th century, deaths from drought declined enormously.



Since 1980, global temperatures have increased by around 0.8°C, but drought deaths have remained very low. This is likely in large part because of economic development.

5.3.4. Heat stress for agricultural workers

Another important impact of climate change is on heat stress for agricultural workers. I discuss this in more detail in Chapter 6.

In short, rising heat stress looks set to reduce labour capacity especially in the tropics, which will be a site of rising population. The effect is summarised in this chart from Buzan and Huber (2020), which assumes that people do not migrate and do not make additional use of air conditioning.

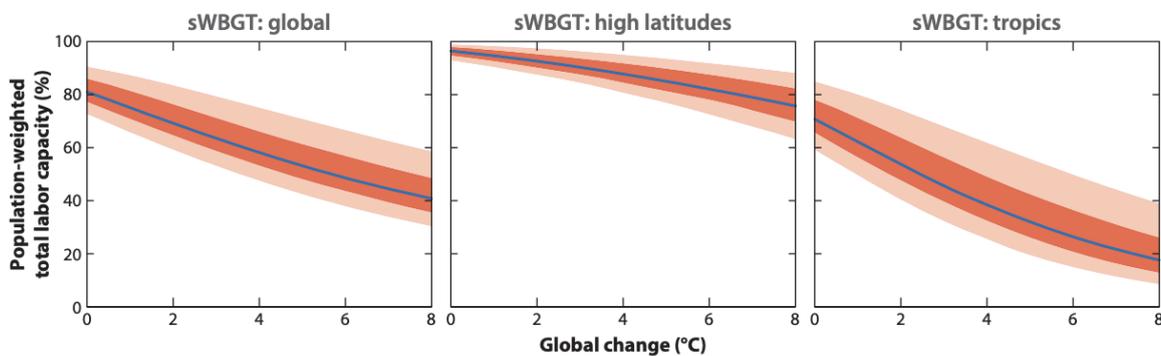


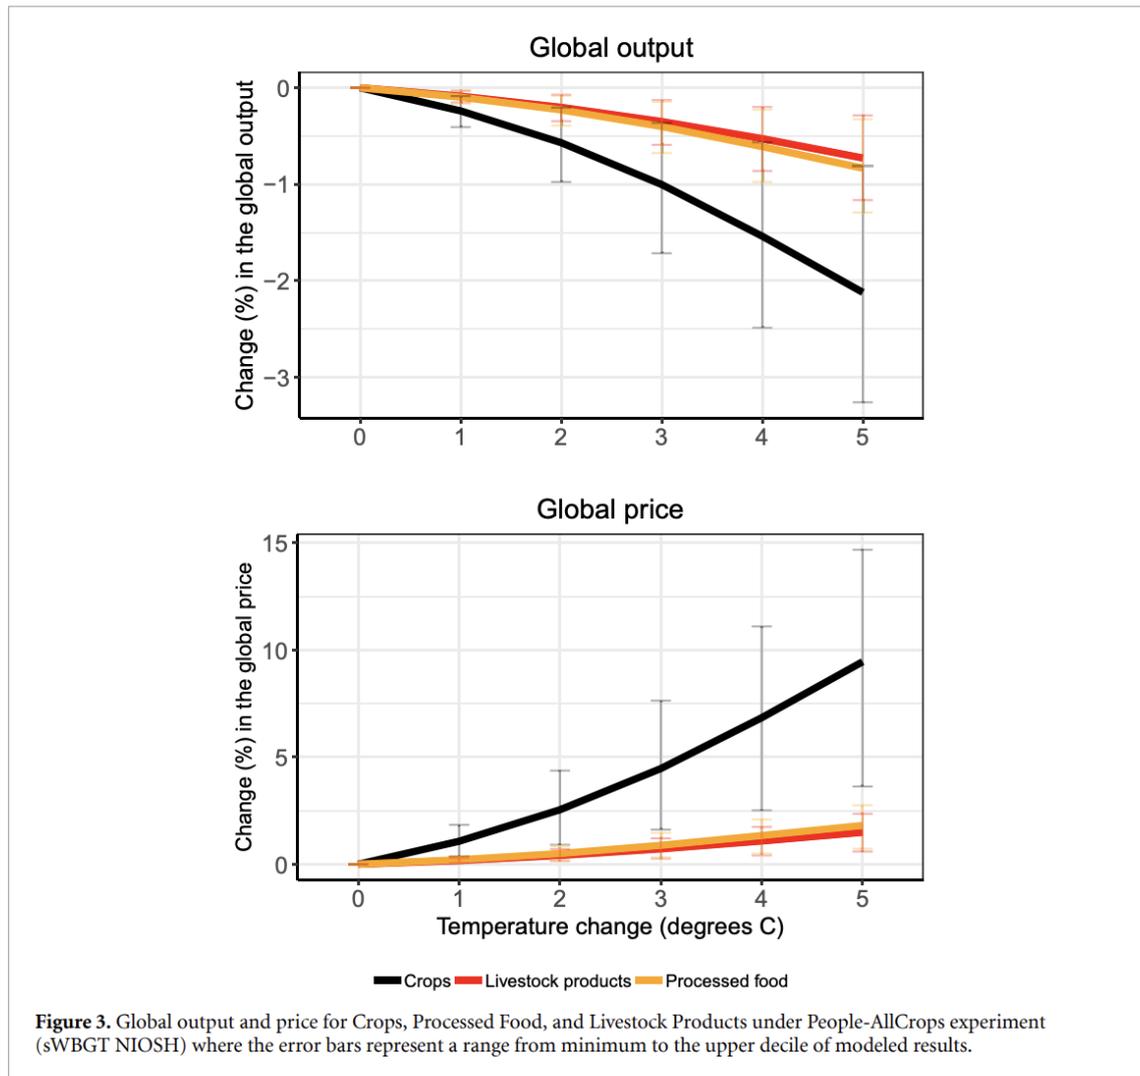
Figure 10

Population weighted total labor capacity. The CMIP5 ensemble is represented by the median (*blue line*), 50% (*red swath*), and 80% (*pink swath*) confidence intervals. The relative impacts on labor are shown at global (57°S to 57°N), high latitude (outside of 30°S to 30°N), and tropic (30°S to 30°N) regions. Abbreviations: CMIP, Coupled Model Intercomparison Project; sWBGT, simplified wet bulb globe temperature.

Source: Jonathan R. Buzan and Matthew Huber, 'Moist Heat Stress on a Hotter Earth', *Annual Review of Earth and Planetary Sciences* 48, no. 1 (2020): 623–55, <https://doi.org/10.1146/annurev-earth-053018-060100>.

There is some scope for countries to adapt to this by using mechanisation, though this response would not be available to very poor countries. Livestock would also be exposed to these rising levels of heat stress, which could be very damaging for animal welfare.

Lima et al (2021) estimate the effect that heat stress will have on agricultural production and prices:



Source: Cicero Z. de Lima et al., 'Heat Stress on Agricultural Workers Exacerbates Crop Impacts of Climate Change', *Environmental Research Letters* 16, no. 4 (March 2021): 044020, <https://doi.org/10.1088/1748-9326/abeb9f>.

5.4. Future agricultural progress

In the last 60 years, crop yields and food production have increased by upwards of 200%. However, as discussed in [this](#) overview by the Breakthrough Institute, there are reasons to think that this progress might slow down in the future. Historical trends and models suggest

that global agricultural yields will increase by 25-150% up to 2050 due to technological progress.²²⁰

Unless there is a huge trend break in agricultural progress, improvements in technology and efficiency look set to outpace the negative effects of climate change. Consequently, we should expect food yields and food production to increase in the future, despite climate change. However, climate change will damage food production, which would be damaging given rising population and rising food demand.

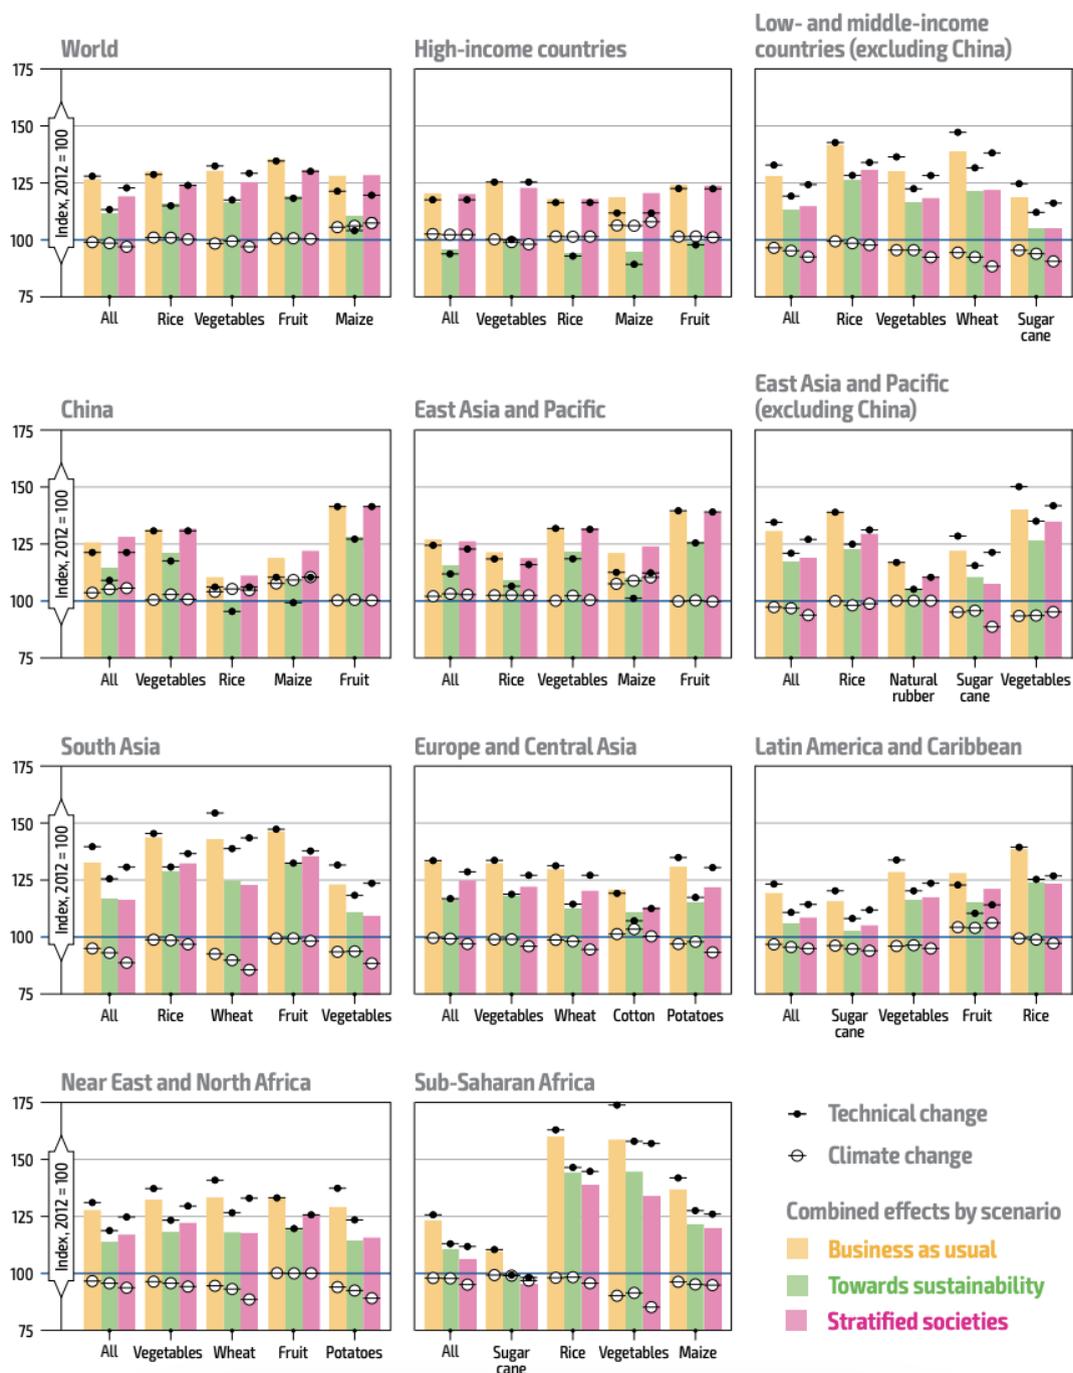
FAO report

A good illustration of this is shown in research from the Food and Agriculture Organisation, which finds that up to 2050 technological change will outpace the effect of climate change up to 2050 in almost all scenarios, in all agricultural systems, and in all regions.

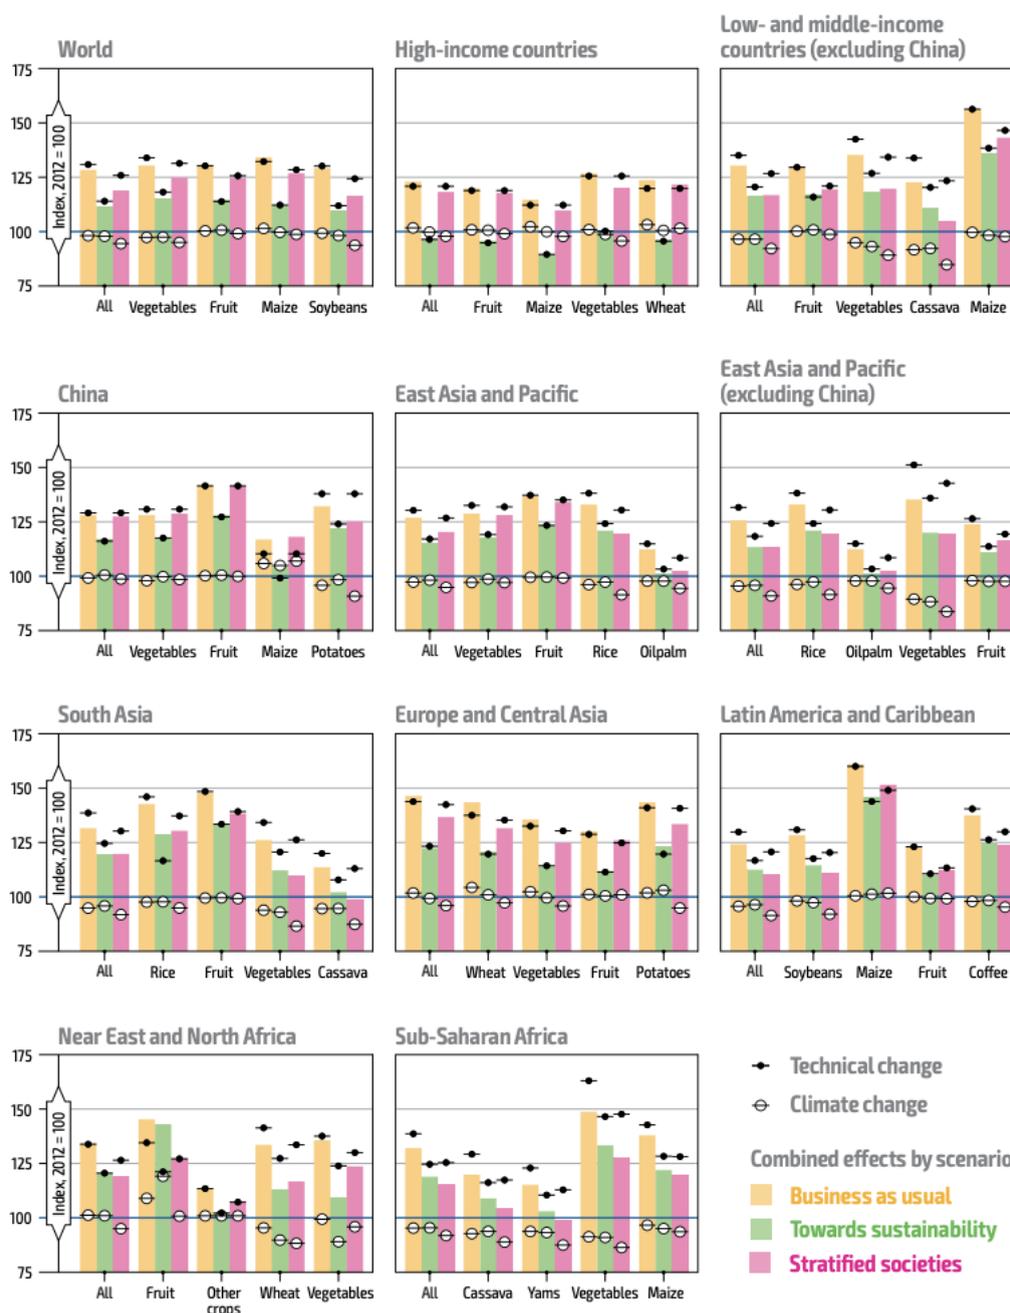
²²⁰ Keith Wiebe et al., 'Climate Change Impacts on Agriculture in 2050 under a Range of Plausible Socioeconomic and Emissions Scenarios', *Environmental Research Letters* 10, no. 8 (August 2015): fig. 3, <https://doi.org/10.1088/1748-9326/10/8/085010>. "We find that yields in these top four crops are increasing at 1.6%, 1.0%, 0.9%, and 1.3% per year, non-compounding rates, respectively, which is less than the 2.4% per year rate required to double global production by 2050. At these rates global production in these crops would increase by ~67%, ~42%, ~38%, and ~55%, respectively, which is far below what is needed to meet projected demands in 2050." Deepak K. Ray et al., 'Yield Trends Are Insufficient to Double Global Crop Production by 2050', *PLOS ONE* 8, no. 6 (19 June 2013): e66428, <https://doi.org/10.1371/journal.pone.0066428>.

Figure 3.9 Yield changes from 2012 to 2050 due to climate change and technical progress

a) Irrigated systems



b) Rainfed systems



Note: Coloured bars indicate price-independent changes in yields attributed to both technical progress and climate change. The white circles indicate changes in yields arising from climate change, while the black barred dots indicate changes arising from technical progress. Climate change impacts are computed based on FAO-IIASA GAEZ v4 (scenario without CO₂ fertilization, median value for five climate models). Changes in yields are shown for the four top commodities, as classified in the FAO GAPS model, in each region, and production system, ranked by value of production in 2012. In this figure, "Citrus" and "Other fruit" are aggregated into "Fruit". "All" refers to the aggregated change in production over the total harvested areas for all crops. Note that the results of research into the impacts of climate change on fruit trees are not conclusive (Ramírez and Kallarackal, 2015).

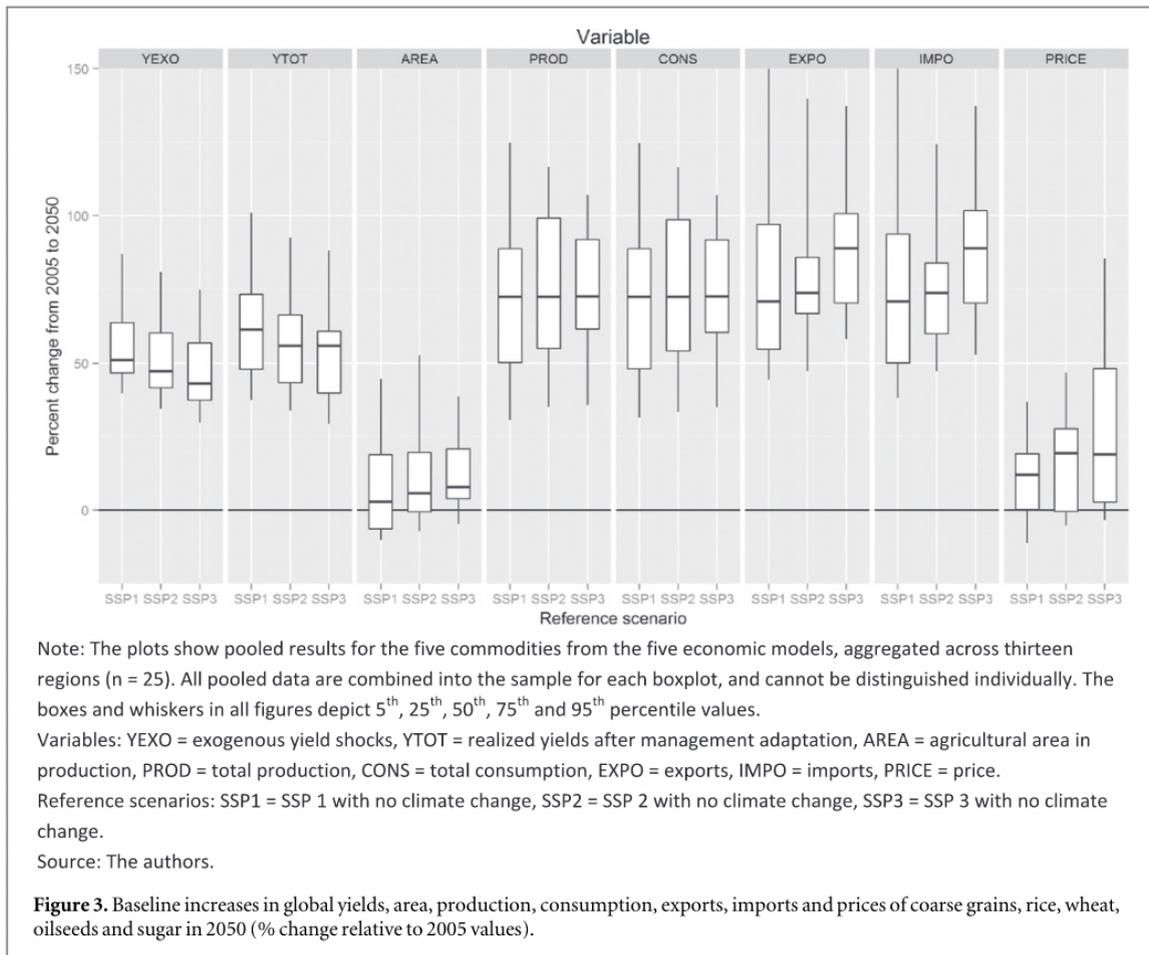
Sources: FAO Global Perspectives Studies, based on FAOSTAT (various years) for historical crop yields and value of production; FAO-IIASA GAEZ v4 for climate change shifters; and FAO expert judgement for technical shifters.

Source: Food and Agriculture Organization, 'The Future of Food and Agriculture. Alternative Pathways to 2050', 2018,

<http://www.fao.org/policy-support/tools-and-publications/resources-details/en/c/1259562/>.

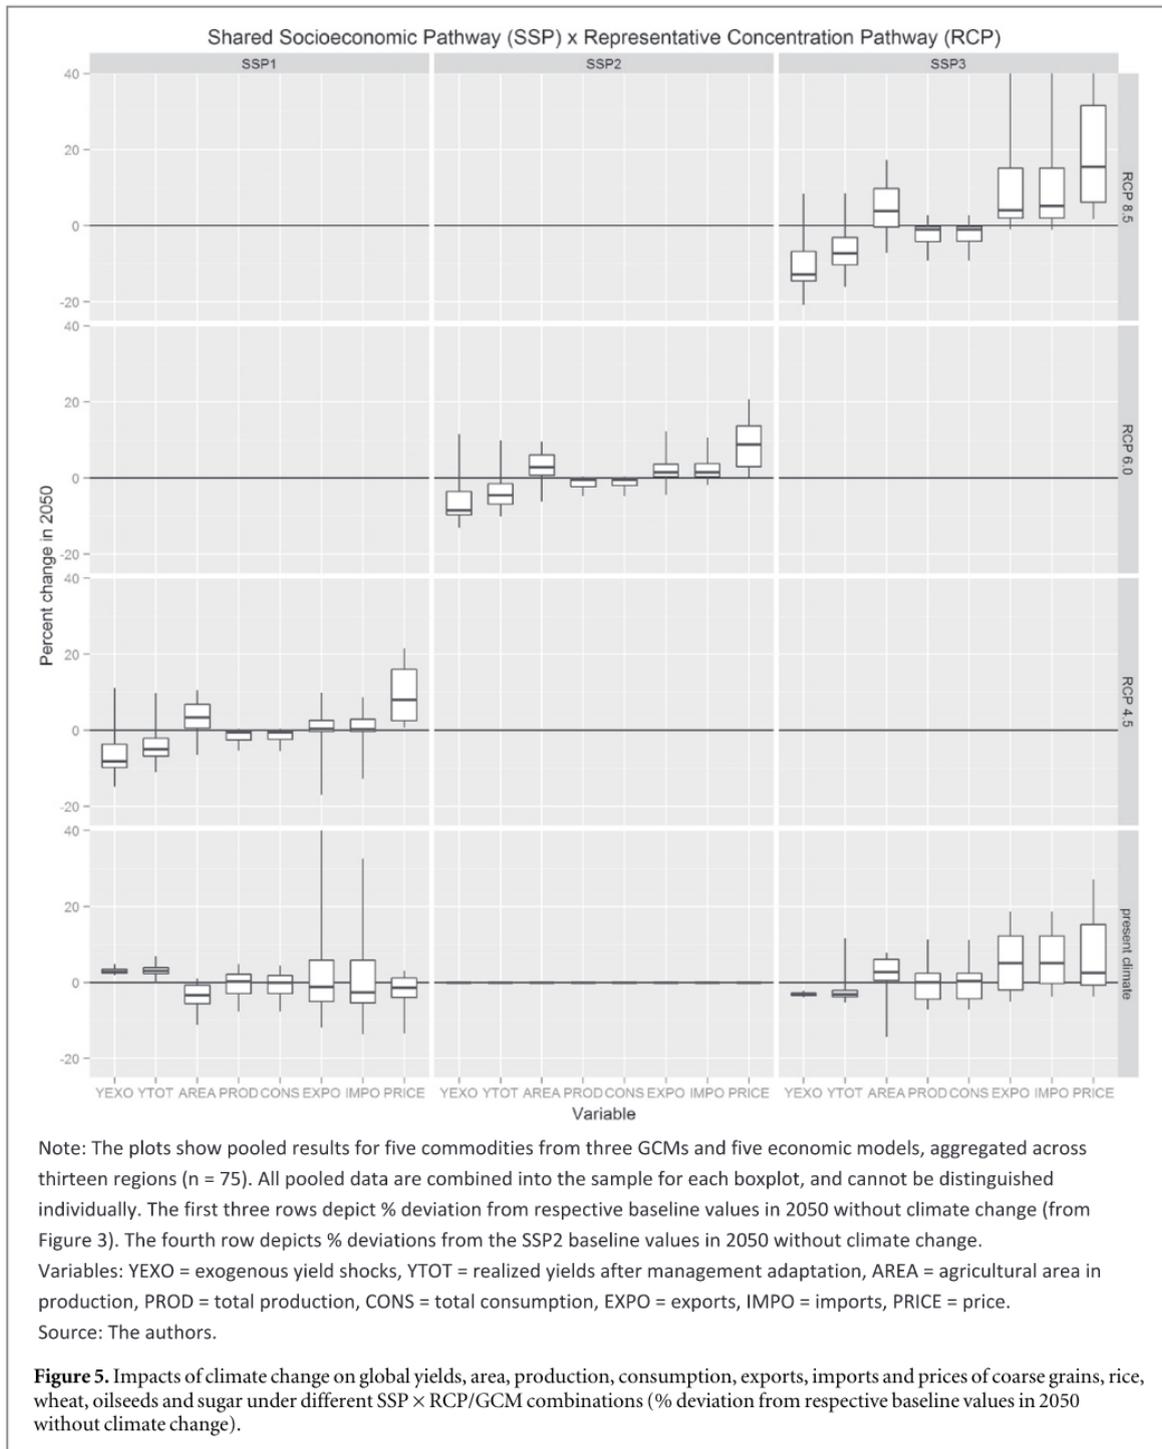
Wiebe et al (2015)

Wiebe et al (2015) project that yields will increase by more than 50% and production will increase by 25% to 125% *relative to today*, on a range of SSPs up to 2050.²²¹



Compared to these improvements, the effect of trend changes in temperature and precipitation are relatively small, though the study does not account for climatic variability. The chart below from Wiebe et al (2015) shows the effect of warming relative to a world without warming in 2050.

²²¹ Keith Wiebe et al., 'Climate Change Impacts on Agriculture in 2050 under a Range of Plausible Socioeconomic and Emissions Scenarios', *Environmental Research Letters* 10, no. 8 (August 2015): 085010, <https://doi.org/10.1088/1748-9326/10/8/085010>.



On RCP8.5, there would be nearly 2.4°C of warming by 2050. As the chart above shows, food production would fall by at most 10% on this scenario *relative to a world without climate change in 2050*. But *relative to today*, production will be higher.

Van Dijk et al (2021)

Van Dijk et al (2021) explores future pathways in food consumption and risk of hunger on different SSPs and RCPs. This chart projects average calories consumed per person. The

panel on the right compares the world without climate change and with climate change, on RCP8.5 by 2050, which implies about 2.4C of warming:

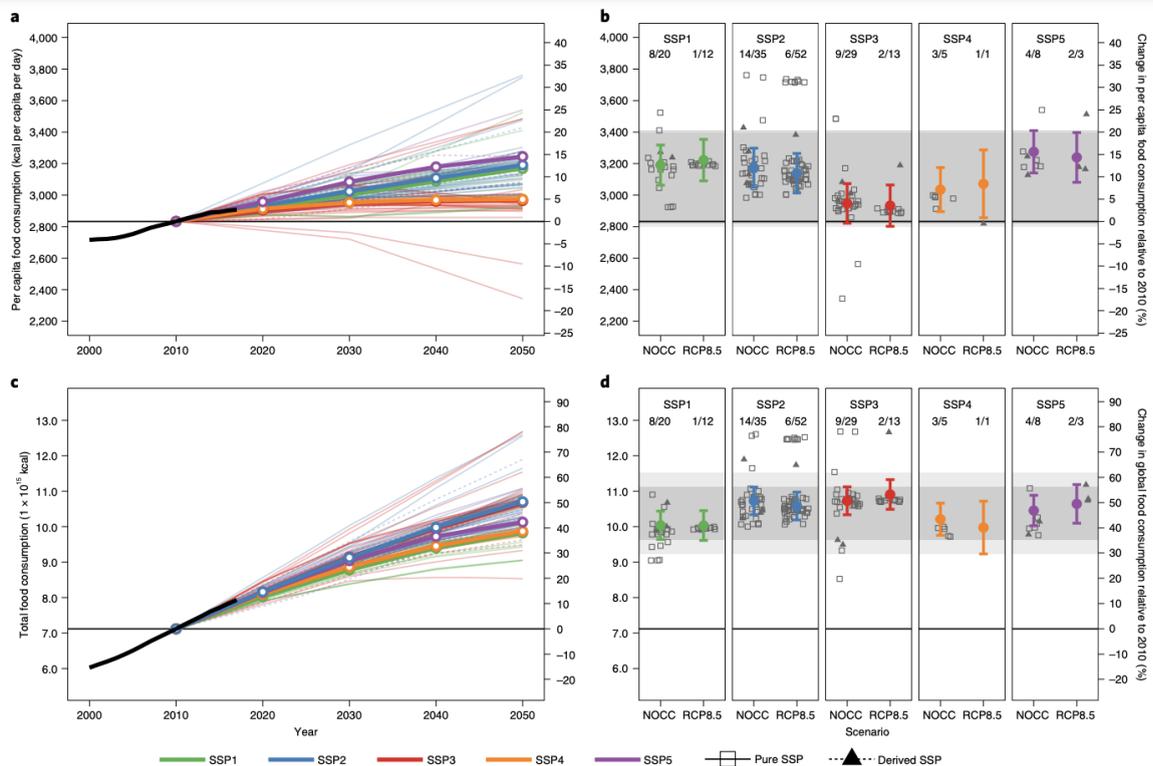


Fig. 3 | Per capita and total food consumption baseline projections for 2010–2050. **a,b**, Per capita food consumption baseline projections. **c,d**, Total food consumption baseline projections. All figures show the level of the selected food security indicator (left y axis) and the percentage increase for the period 2010–2050 (right y axis). Panels **a** and **c** show the baseline projections for the SSPs under NOCC (thin coloured lines), the average for each SSP (thick coloured lines with circles) and the three-year average historical trend (thick black line). Panels **b** and **d** present point estimates and 95% confidence intervals taken from the meta-regression as well as all observations in 2050, comparing NOCC projections with projections based on the most extreme climate scenario (RCP8.5). The numbers at the top refer to the number of studies/number of projections in the figure. The dark and light grey shaded areas demarcate the plausible range of baseline projections using the 95% confidence interval across all NOCC SSP and all RCP SSP projections, respectively (Extended Data Fig. 4). Pure SSPs are projections that take their assumptions from the SSPs, where relevant combined with RCP-based climate impact scenarios. Derived SSPs are projections that belong to the same SSP and RCP scenario families but use somewhat different assumptions. The historical data are from FAO⁷⁰. The projections are from the Global Food Security Projections Database.

Source: Michiel van Dijk et al., 'A Meta-Analysis of Projected Global Food Demand and Population at Risk of Hunger for the Period 2010–2050', *Nature Food* 2, no. 7 (July 2021): 494–501, <https://doi.org/10.1038/s43016-021-00322-9>.

As this shows, per capita calorie consumption looks set to increase with 2.4°C of warming. On most SSPs, the fraction of the population at risk of hunger will decrease in the future, even with 3°C of warming, as shown in the right pane below.

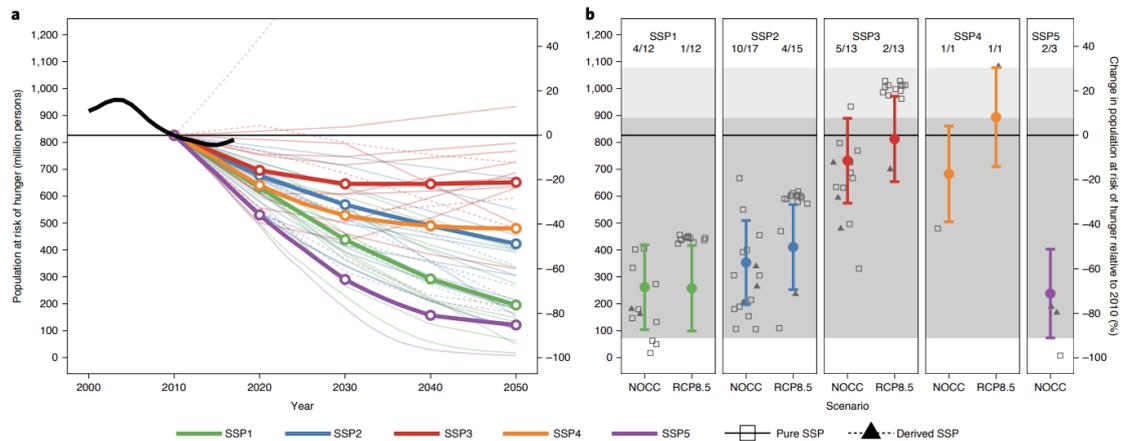


Fig. 4 | Population at risk of hunger baseline projections for 2010-2050. **a**, Baseline projections for the SSPs under NOCC (thin coloured lines), the average for each SSP (thick coloured lines with circles) and the three-year average historical trend (thick black line). **b**, Point estimates and 95% confidence intervals taken from the meta-regression as well as all observations in 2050, comparing NOCC projections with projections based on the most extreme climate scenario (RCP8.5). See Fig. 3 for a detailed explanation of the figure elements. The dark and light grey shaded areas demarcate the plausible range of baseline projections using the 95% confidence interval across all NOCC SSP and all RCP SSP projections, respectively (Extended Data Fig. 5). The projections from Dawson et al.⁷² (the blue dashed line heading upwards in **a**) are considered outliers and are therefore excluded from the meta-regression. The historical data are from FAO⁷⁰. The projections are from the Global Food Security Projections Database.

This illustrates that socioeconomic development and economic growth will be the most important determinant of the number of people at risk of hunger in the future. The difference between the high growth future - SSP5 - and the low growth futures - SSP3 and SSP4 - are much larger than the projected effect of climate change.

5.5. Overall judgement on the effects of climate change on agriculture

Climate change will damage agricultural output via a variety of mechanisms. The effects will be worst for people in the tropics who have contributed the least to climate change and are least able to adapt. These will also be the sites of the greatest future population growth. This gives us a strong reason to reduce emissions and to encourage socioeconomic development in the affected regions.

However, the existing evidence suggests that, with warming of 4°C, total global food production will very probably be higher than today. Even though food demand will rise this century, food consumption per person will also very probably be higher than today. There is some evidence that synchronised food production declines for some major food crops will increase, but these shocks will occur against a baseline in which food yields and production are much higher.

General equilibrium effects via the price system will also attenuate some of the humanitarian costs of higher food prices. The world's food producing regions are geographically spread out and climatically diverse, and there are enormous internalised market incentives to produce food and to respond to the changing climate through crop switching, crop migration, genetic modification, changes in cultivation practices, and so on. The overall effects of food production on human welfare can be captured by economic models, which I discuss in Chapter 10.

The effects of warming of 5°C and above on food production are not well-studied. The limited modelling evidence that does exist suggests that yield losses would be around 20% to 50% for 6°C of warming. Given how long it would take to reach 6°C, overall yields would probably still be higher, assuming that society has not collapsed or stagnated for some other reason. The effects of heat stress on agricultural workers is small relative to direct climatic effects on crops.

I will now discuss the risks of extreme warming to agriculture in more depth.

5.6. Extreme warming and agriculture

What level of warming would threaten the global viability of agriculture? I have been unable to find any studies that explore this question systematically, so will attempt my own answer.

5.6.1. Lethal limits

The majority of plant species on Earth use [C₃ photosynthesis](#), in which the first carbon compound produced by photosynthesis contains three carbon atoms (hence the name). In C₄ photosynthesis, plants produce a compound containing four carbon atoms. 85% of plant species are C₃ plants including important sources of calories such as wheat, rice, barley, oats, cowpeas, cassava, soybeans.

Once temperatures pass 35°C, for C₃ plants photorespiration starts to predominate over photosynthesis.

Figure 1. Observed (FAO 2003) grain yields of wheat for selected countries in Europe.

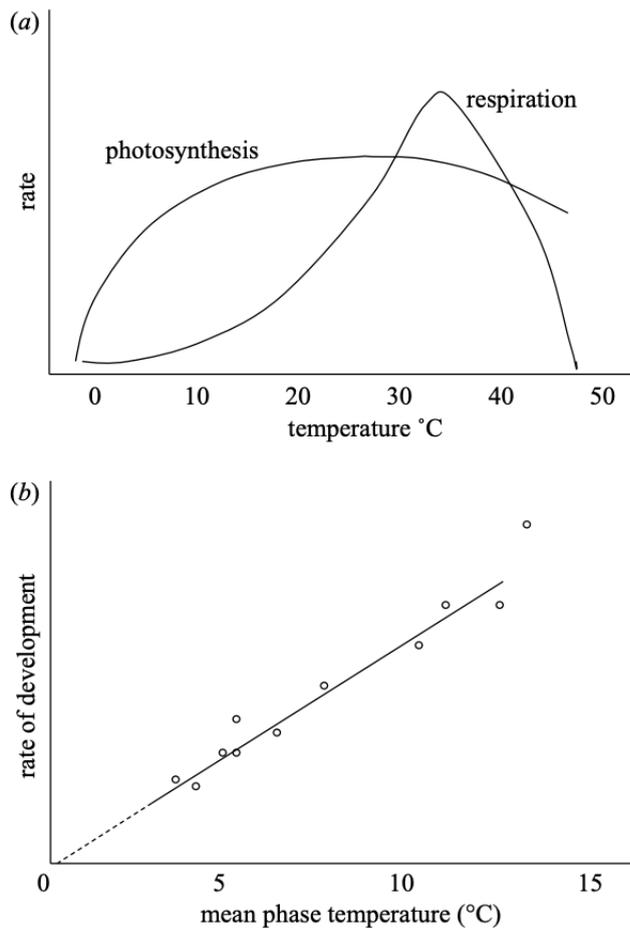


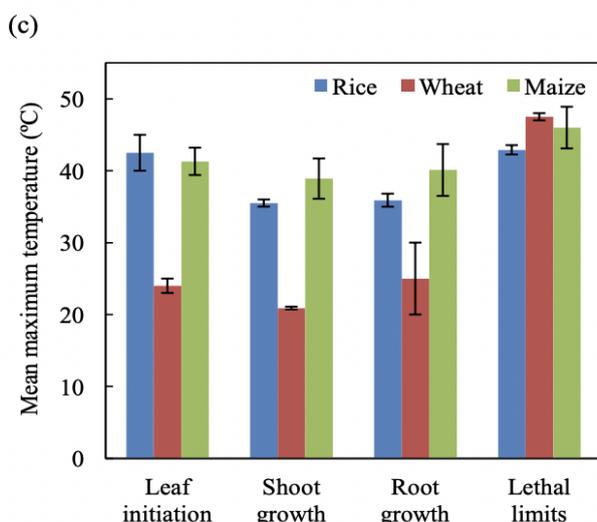
Figure 2. Changes in the rate of (a) C_3 photosynthesis and respiration and (b) rate of crop development as a function of temperature.

Source: John R Porter and Mikhail A Semenov, 'Crop Responses to Climatic Variation', *Philosophical Transactions of the Royal Society B: Biological Sciences* 360, no. 1463 (29 November 2005): 2021–35, <https://doi.org/10.1098/rstb.2005.1752>.

Few C_3 plants can survive temperatures persistently above 40°C.²²² The figure below shows a range of thresholds for wheat, rice and maize:

²²² "For C_3 plants, photorespiration predominates over photosynthesis at temperatures in excess of 35°C (13), and few plants can survive temperatures persistently above 40°C (14)." Yadong Sun et al., 'Lethally Hot Temperatures During the Early Triassic Greenhouse', *Science* 338, no. 6105 (19 October 2012): 366–70, <https://doi.org/10.1126/science.1224126>.

Figure 3: Rice, wheat and maize - Mean maximum temperature for leaf initiation, shoot growth, root growth and lethality.¹⁰



Source: David King et al., 'Climate Change—a Risk Assessment' (Centre for Science Policy, University of Cambridge, 2015), p. 68. www.csap.cam.ac.uk/projects/climate-change-risk-assessment²²³

'Maximum temperature' here means that the relevant process will be put on hold rather than permanently stopped.²²⁴ For instance, while temperatures are beyond 28°C, leaf initiation will not occur in wheat, but once temperatures drop below 28°C, leaf initiation can occur. Soybean is subject to similar lethal limits.²²⁵ Lethal limits would have to be passed for around 1-5 days to kill the plant.²²⁶

The lethal limit for each of these crops is 42-47°C. I will assume in what follows that 40°C is the lethal limit.

At what level of warming would plants pass lethal limits?

How high would global warming have to get for temperatures to be above 40°C for a sustained period in key food producing regions, with effects potentially sufficient to kill more than 50% of the global population?

I am not aware of any studies that try to answer this question. So, I will attempt a rough answer myself. I will try to determine when temperatures would pass lethal limits for major food crops in food production in five key areas: North America, Europe, China, India and Russia. I include Russia because a lot of its frozen land would be freed up for agriculture, on

²²³ See also Porter and Semenov, 'Crop Responses to Climatic Variation', Table 2.

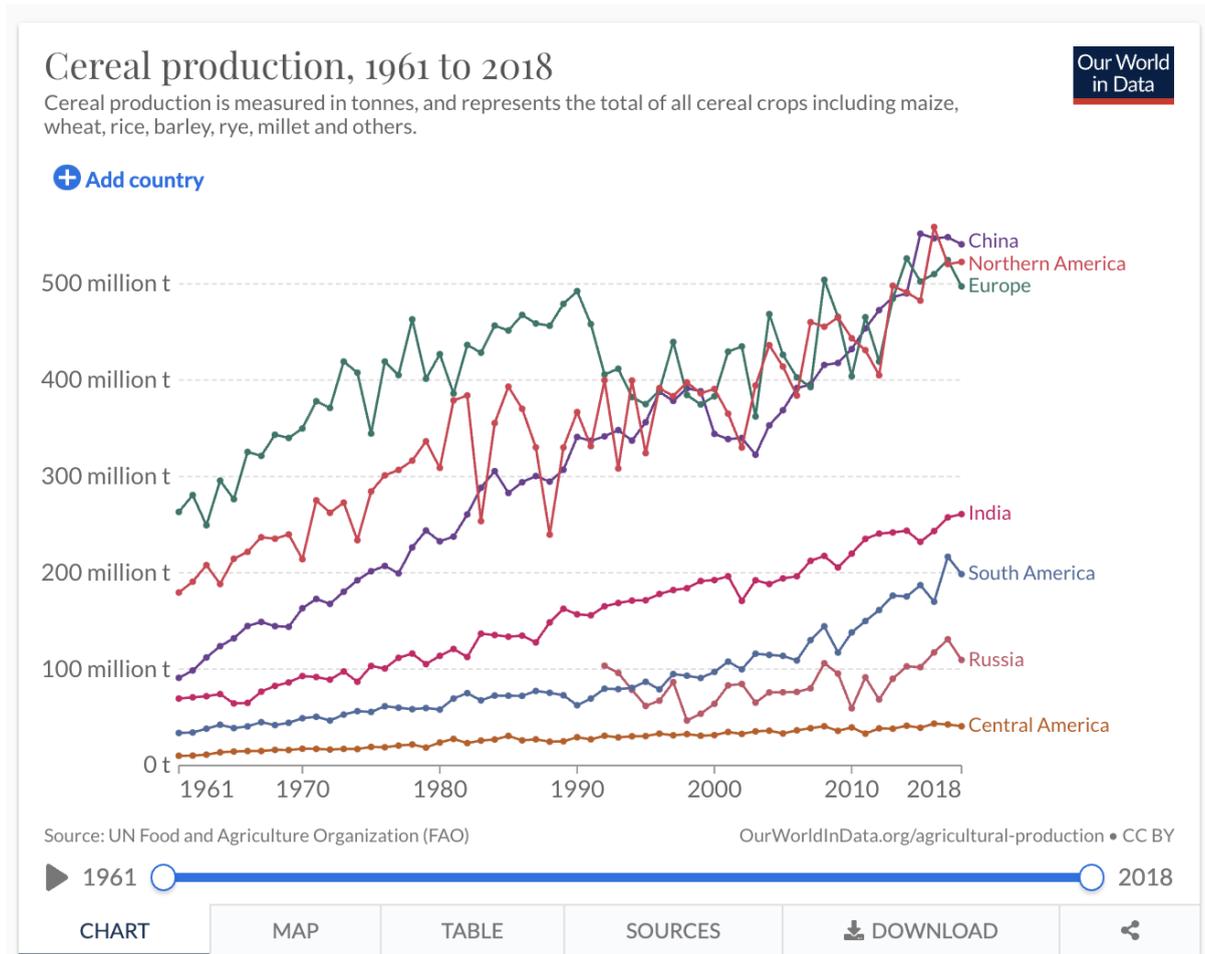
²²⁴ Mikhail Semenov, personal correspondence, 12 Aug 2021.

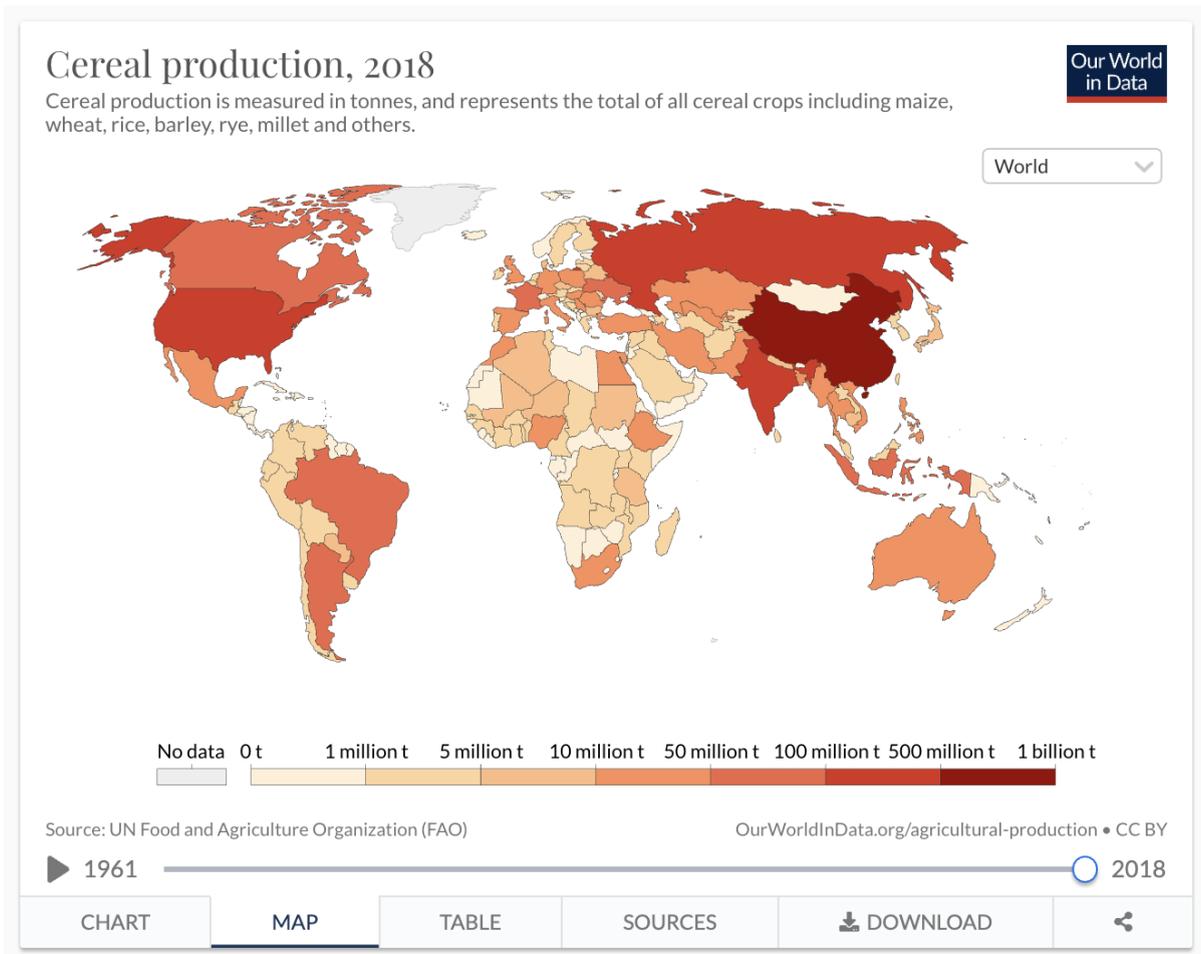
²²⁵ Wolfram Schlenker and Michael J. Roberts, 'Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change', *Proceedings of the National Academy of Sciences* 106, no. 37 (15 September 2009): fig. 1, <https://doi.org/10.1073/pnas.0906865106>.

²²⁶ "Figure 5 shows the influence of lethal temperatures on simulated yield for thresholds of 40, 45, and 50 °C and their exceedence for 1, 3, and 5 consecutive days in order to lead to plant death." Ann-Kristin Koehler et al., 'Influences of Increasing Temperature on Indian Wheat: Quantifying Limits to Predictability', *Environmental Research Letters* 8, no. 3 (August 2013): 034016, <https://doi.org/10.1088/1748-9326/8/3/034016>.

extreme warming scenarios. Together, these regions today account for 70% of global cereal production. North America, China and Europe alone account for 57% of global cereal production.

I will consider warm and cold locations in each of these regions in order to represent a reasonable spread of climatic conditions in the relevant regions.



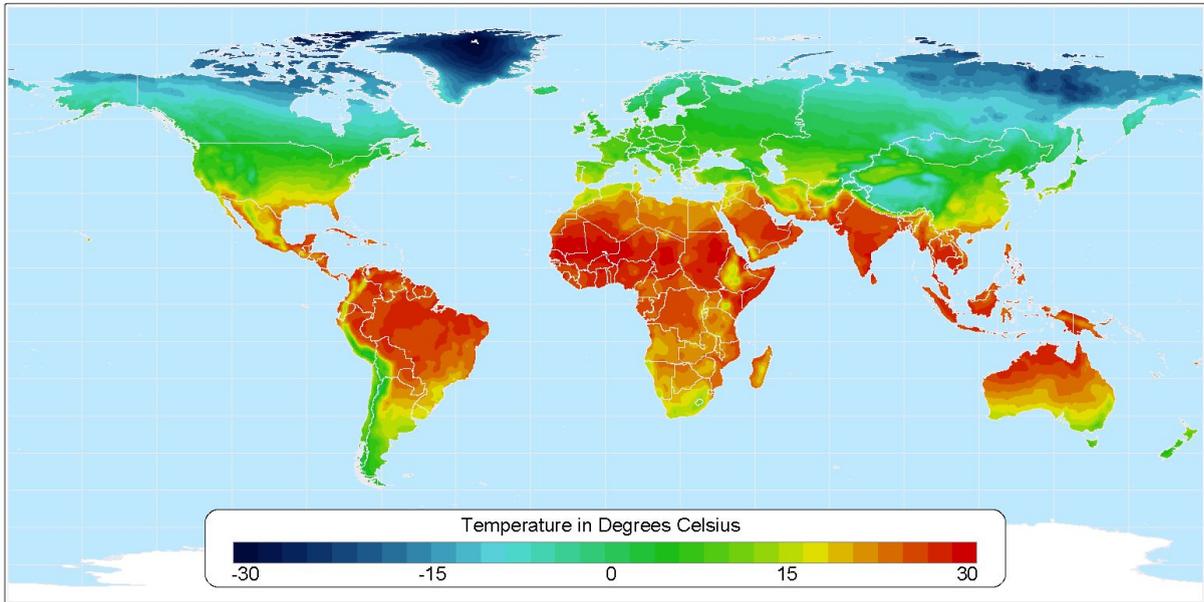


My sole focus in this section is on when different crops start to pass lethal limits. So, it is important to stress that there are many things my approach will leave out:

- Heat stress for agricultural workers
- Changes in soil moisture and drought
- Changes in pest distributions

I also do not consider how appropriate the other determinants of agricultural production, such as soil quality and type, might be in the regions considered. However, I will try to compare specific regions where agriculture is currently carried out at large scale. The exception to this is cold regions in Russia, which I investigate because they are relevant to extreme future warming. In general, my answer will be quite rough and back of the envelope but, I hope, better than nothing.

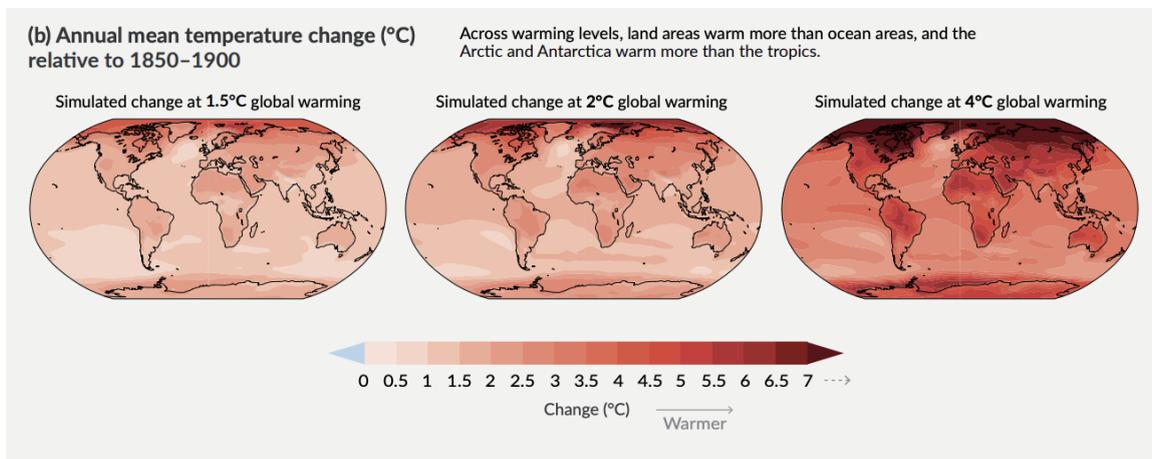
The map below shows annual average temperatures across the world.



Data taken from: CRU 0.5 Degree Dataset (New, et al.)

Atlas of the Biosphere
 Center for Sustainability and the Global Environment
 University of Wisconsin - Madison

Global warming will not be uniform across the world. Warming is expected to be higher at high latitudes, and is generally larger over land than the oceans.



Source: IPCC, *Climate Change 2021: The Physical Science Basis, Sixth Assessment Report* (UNFCCC, 2021), Summary for Policymakers, SPM.5.

However, annual average temperature masks significant seasonal and diurnal variation, which has an important bearing on the prospects for agriculture.

The key risk to food crops from warming is that at some point in the period from planting to harvest, temperatures pass lethal limits. It is therefore important to consider the planting seasons for the major food crops

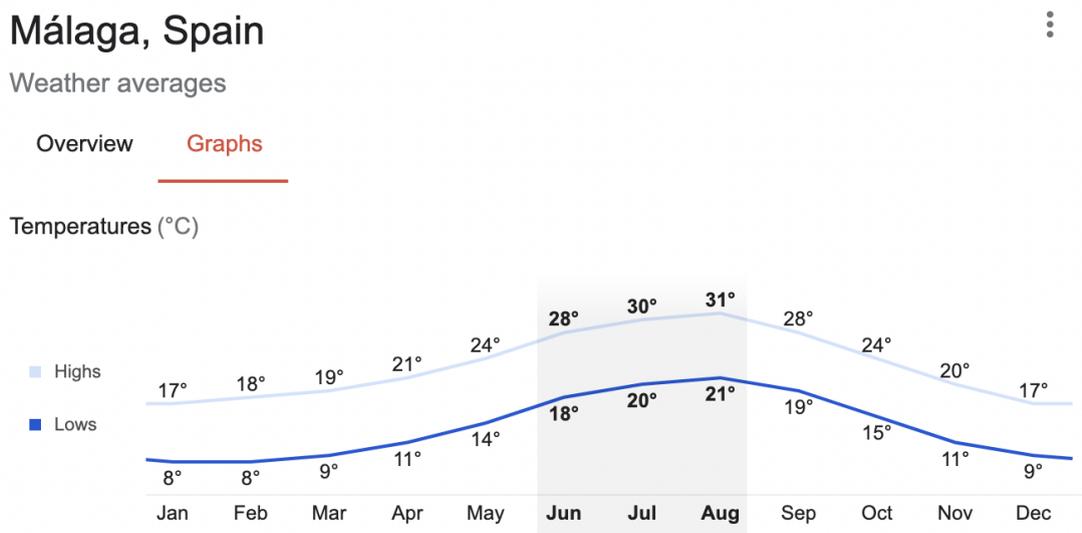
- Wheat ([source](#))

- Winter wheat: planted in the autumn, 7-8 months to reach full maturity, harvested in spring or early summer.
- Spring wheat: planted in spring, 4 months to maturity, harvest in summer or early autumn.
- Rice ([source](#))
 - Plant in late winter/early summer, 4 months to reach maturity, harvest in summer to early winter.
- Maize ([source](#))
 - Plant in spring/summer, 4 months to reach maturity, harvest in autumn.
- Soy ([source](#))
 - Late spring/early summer, 2 months to reach maturity, harvest in summer/autumn.

Europe

To capture a reasonable range across weather in Europe, I will compare Málaga in southern Spain (a hot region) to Edinburgh (a cold region).

The charts below shows changes in monthly temperatures in Málaga



The temperature highs shown are mean daily highs in a given month. So, the highs would mask some within-month variation. Data for [August 2021](#) suggests that temperatures might vary by around 6°C either side in a given month. So, we can assume that the true temperature high in the hottest months is 6°C higher.

If there was local warming of 6°C, then summer temperatures would approach lethal limits in three of the summer months. However, the growing season for the major food crops is 8 months or less, so agriculture would still be viable even if temperatures were this high.

To destroy the growing season for most major food crops, lethal limits would have to be passed for more than 8 months in a year. This would happen once local warming reached 15°C. It is unclear what level of global average surface warming would be required to

produce 15°C of warming in Málaga, but, based on the IPCC map of future regional temperatures, ~10°C of global average surface warming is a reasonable bet.

Edinburgh, UK



Weather averages

Overview **Graphs**

Temperatures (°C)



- Lethal limits in summer: 17°C local warming²²⁷
- Wipe out 8 months of the growing season: 24°C.
 - Equates to ~20°C of global average warming.

North America

For North America, I will compare Houston, Texas (a hot region) with Calgary, Alberta (a cold region).

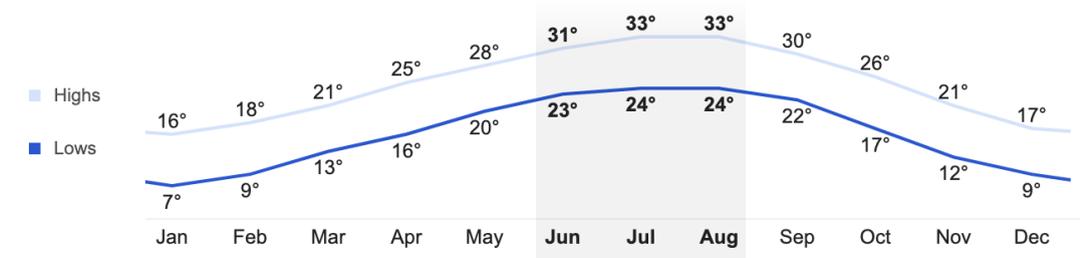
Houston, TX, USA



Weather averages

Overview **Graphs**

Temperatures (°C)



- Lethal limits in summer: 5°C local warming
- Wipe out 8 months of the growing season: 15°C.
 - Equates to ~10°C of global average warming.

²²⁷ This is using the same method used for Málaga, described above.

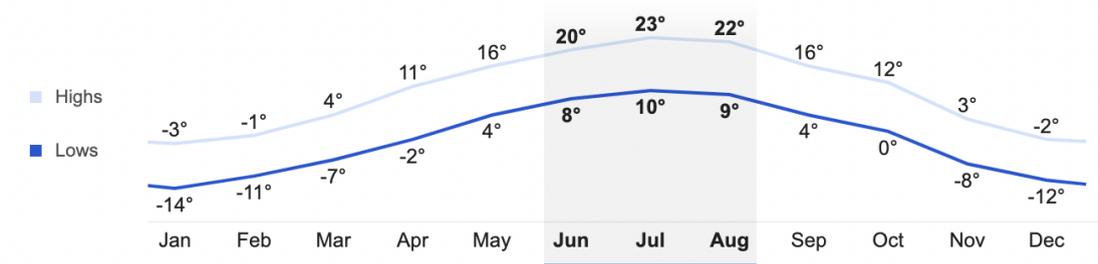
Calgary, AB, Canada



Weather averages

Overview Graphs

Temperatures (°C)



- Lethal limits in summer: 14°C local warming
- Wipe out 8 months of the growing season: 31°C.
 - Equates to ~15°C of global average warming because warming is so much higher at high latitudes.

China

For China, I will compare Guangzhou (a hot region) with Shenyang (a colder region).

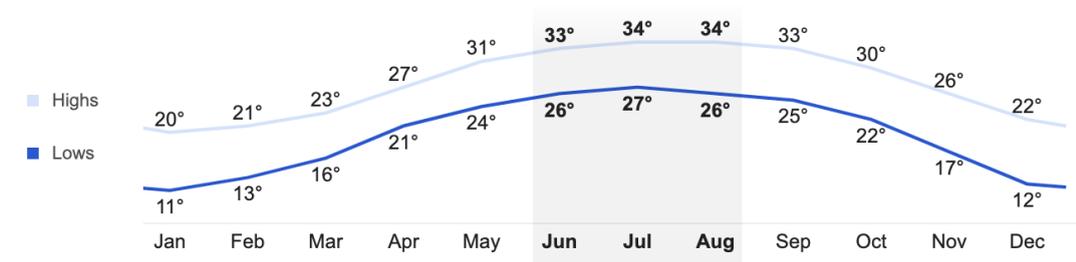
Guangzhou, Guangdong Province, China



Weather averages

Overview Graphs

Temperatures (°C)



- Lethal limits in summer: 3°C local warming
- Wipe out 8 months of the growing season: 11°C.
 - Equates to ~11°C of global average warming

Shenyang, Liaoning, China



Weather averages

Overview **Graphs**

Temperatures (°C)



- Lethal limits in summer: 10°C local warming
- Wipe out 8 months of the growing season: 31°C
 - Equates to ~31°C of global average warming

Russia

For Russia, I will compare Moscow (a fairly climatically typical city) to Yakutsk (the coldest city on Earth)

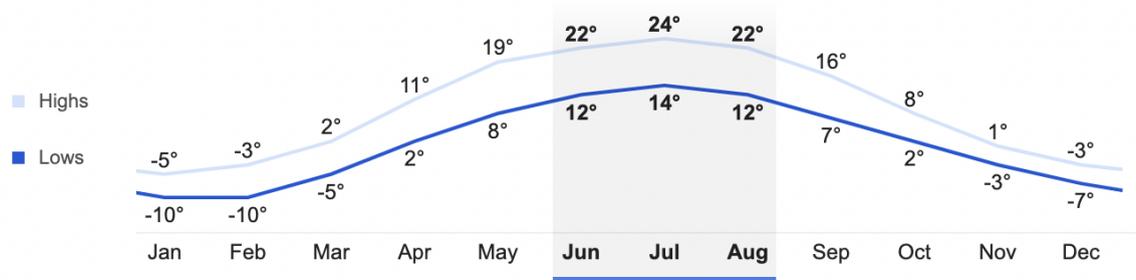
Moscow, Russia



Weather averages

Overview **Graphs**

Temperatures (°C)



- Lethal limits in summer: 13°C local warming
- Wipe out 8 months of the growing season: 34°C.
 - Equates to ~20°C of global average warming

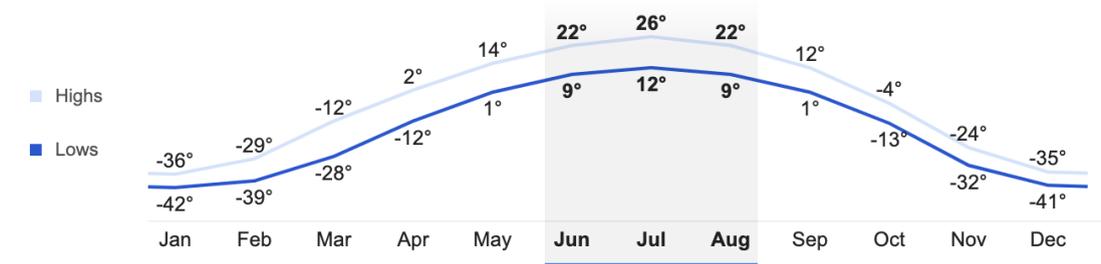
Yakutsk, Sakha Republic, Russia



Weather averages

Overview **Graphs**

Temperatures (°C)



- Lethal limits in summer: 13°C local warming
- Wipe out 8 months of the growing season: 59°C.
 - Equates to ~30°C of global average warming

India

There is less variation in temperatures across India than in the regions I have considered so far. To illustrate the variation, I choose Kochi, a city in the south of India, and New Delhi in the north.

Kochi, Kerala, India



Weather averages

Overview **Graphs**

Temperatures (°C)



- Wipe out 8 months of the growing season: 9°C.
 - Equates to ~9°C of global average warming.

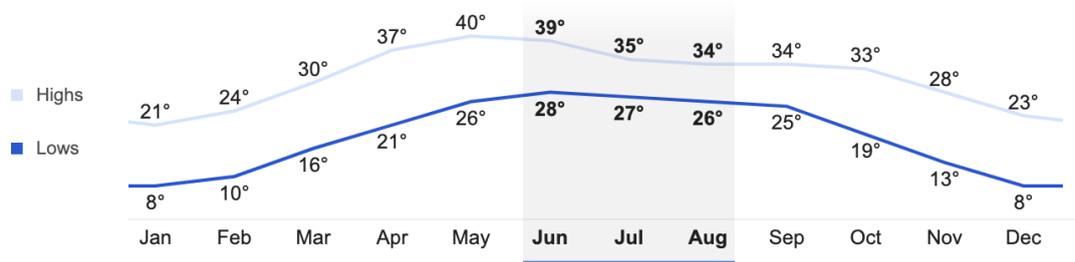
New Delhi, Delhi, India



Weather averages

Overview **Graphs**

Temperatures (°C)



- Wipe out 8 months of the growing season: 10°C.
 - Equates to ~10°C of global average warming.

These findings are summarised in the table below

Region	City	Global average warming to destroy 8 months of growing season
Europe	Málaga	10°C
	Edinburgh	20°C
North America	Houston	10°C
	Calgary	15°C
China	Guangzhou	11°C
	Shenyang	31°C
Russia	Moscow	20°C
	Yakutsk	30°C
India	Kochi	9°C
	New Delhi	10°C

To repeat, this exercise only assesses thermal limits to plants, and not other determinants of food production.

This analysis suggests that more than 10°C of global average surface warming would destroy agriculture in India and regions with similar climates, as well as warm regions in China, Europe and North America. Warming of 15°C would destroy agriculture in North America, while warming of around 30°C would destroy agriculture in China and Russia.

These estimates are uncertain, but I wouldn't expect them to be wrong by more than 6°C either side.

5.6.2. Extreme heat stress and agricultural labour

In the previous section I discussed the findings of de Lima et al (2021) which found that for 5°C of warming, agricultural production would decline by 1-3% due to heat stress. I am not aware of any studies quantifying the effect of more than 5°C of warming.

The findings of Buzan and Huber (2020) on labour supply are shown below

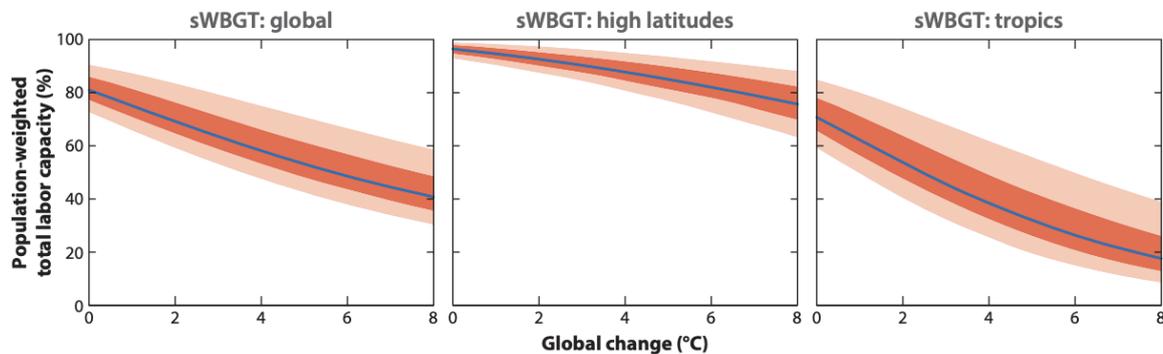


Figure 10

Population weighted total labor capacity. The CMIP5 ensemble is represented by the median (*blue line*), 50% (*red swath*), and 80% (*pink swath*) confidence intervals. The relative impacts on labor are shown at global (57°S to 57°N), high latitude (outside of 30°S to 30°N), and tropic (30°S to 30°N) regions. Abbreviations: CMIP, Coupled Model Intercomparison Project; sWBGT, simplified wet bulb globe temperature.

Source: Buzan and Huber, 'Moist Heat Stress on a Hotter Earth'.

Buzan and Huber (2020) find that, assuming people do not move and do not take adaptation measures, agricultural labour would be all-but impossible in the tropics for 8°C of warming, and at the global level, population-weighted labour capacity would decline to 50% of its potential. Given the pessimistic assumptions about response and adaptation, this estimate is plausibly on the high side, but the scope for adaptation in agriculture is also limited.

5.6.3. Summary of extreme warming and agriculture

I argued in Chapter 2 that on a worst-case scenario in which we burn all the fossil fuels, warming would most likely reach 7°C, and there is a 1 in 6 chance of more than 9.5°C of warming. The evidence we have suggests that even this level of warming would not come close to destroying global agriculture, though it would have disastrous effects. Agricultural land would be freed up in Russia and Canada, which would offset some of the costs in other countries.

In general, the persistence of agriculture relies on:

- Sunlight
- Enough rain and water resources in enough regions of the world
- Sufficient CO₂ levels
- Temperatures not falling below freezing in enough regions of the world, which destroys the growing season

- Temperatures not rising above lethal limits for major food crops in enough regions of the world

It is easy to see how an event like nuclear winter could threaten the global viability of agriculture. A nuclear winter would block out the Sun, and would cause temperatures to drop below freezing in many regions, which would destroy the growing season. In contrast, it is difficult to see how plausible levels of global warming could do comparable damage.

- There would still be sunlight and enough CO₂ to allow photosynthesis.
- Lethal limits for the major food crops are a very long way away.
- No climate models project that rain will stop completely due to climate change. In fact, global average precipitation will increase.²²⁸ Some regions would get wetter and some would get drier. There would still be enough water resources in some regions to maintain agriculture.

Although climate change will be damaging to agriculture, it is difficult to come up with realistic scenarios in which food production would decline by more than half due to climate change.

One important possible caveat to this is the risk of tipping points, which I discuss in Chapter 8.

5.7. Ecosystem collapse and threats to agriculture

One possibility is that climate change could lead to ecosystem collapse which would undermine global agriculture. The causal chain would look something like this:

Global warming + habitat loss + human predation + pollution => global species loss
=> global ecosystem collapse => global agricultural catastrophe

For example, Steffen et al (2015) propose that reducing biodiversity too far below pre-industrial levels would consist in crossing a 'planetary boundary'. According to Steffen et al (2015), crossing these boundaries would greatly increase the risk of driving "the Earth system to a much less hospitable state".²²⁹

This argument focuses on the *instrumental* benefits of ecosystems. Ecosystems may also have intrinsic value, but my focus here is the benefits they provide to human civilisation.

5.7.1. Trends in biodiversity

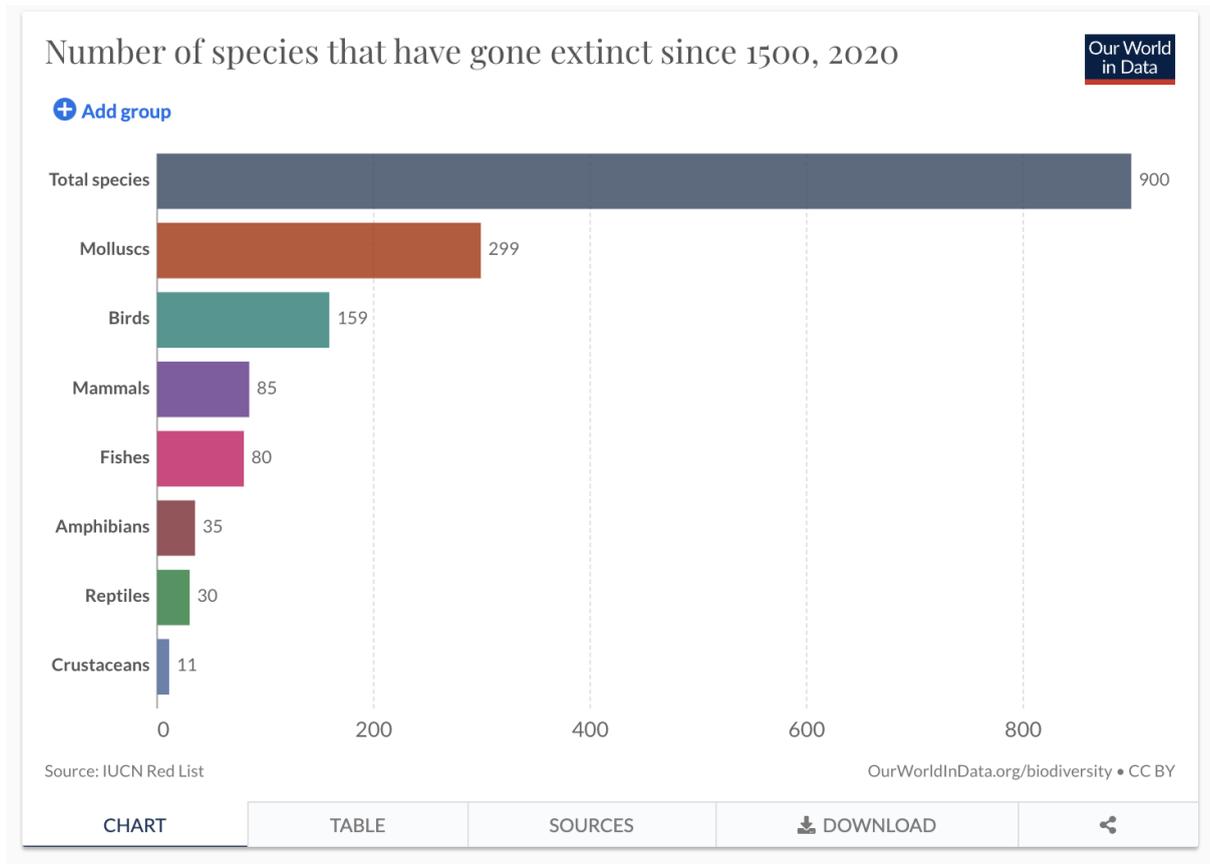
Our World in Data recently released several entries about [biodiversity](#), which this section relies on.

²²⁸ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, p. 19.

²²⁹ Will Steffen et al., 'Planetary Boundaries: Guiding Human Development on a Changing Planet', *Science* 347, no. 6223 (13 February 2015): p. 737, <https://doi.org/10.1126/science.1259855>.

Species extinctions

Since 1500, 900 species have been recorded as extinct.



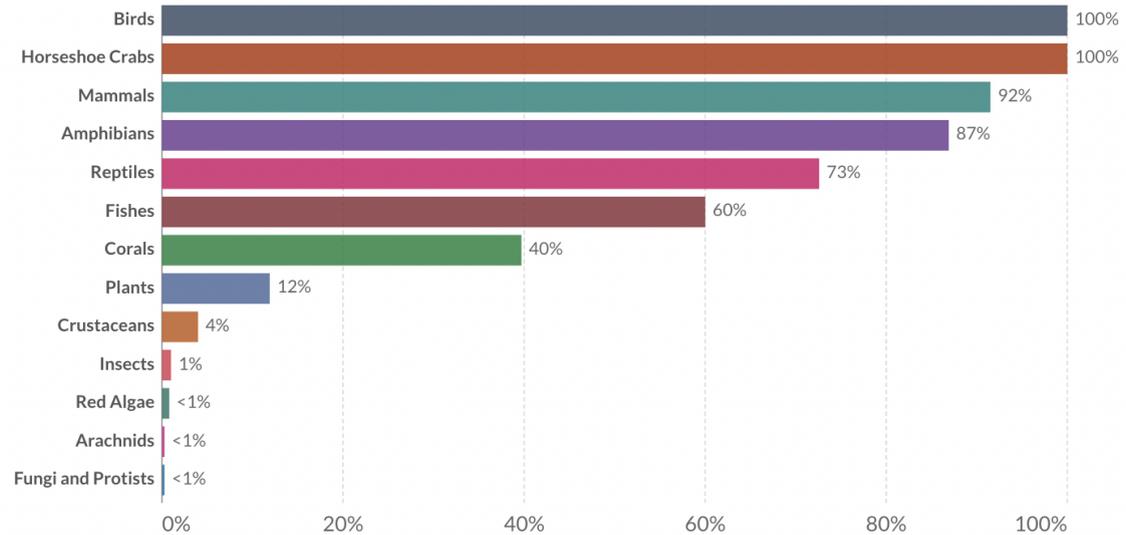
However, the share of species that have been evaluated for extinction varies across taxonomic groups

Share of described species that have been evaluated for their extinction risk, 2020

Our World
in Data

In many taxonomic groups, very few described species have been evaluated for their extinction risk level. This means the estimated number of species at risk of extinction in these groups is likely to be a significant undercount.

[+ Add taxonomic group](#)



Source: IUCN Red List

OurWorldInData.org/biodiversity • CC BY

CHART

TABLE

SOURCES

DOWNLOAD



This means that the estimate of the number of extinct and threatened species is an underestimate.

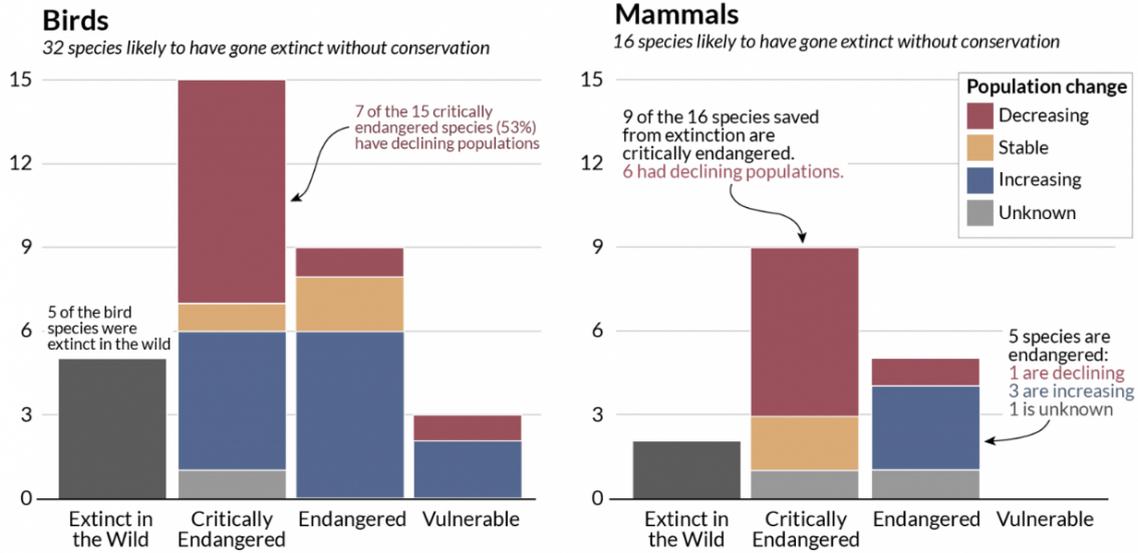
Threatened species are the sum of the following three categories

- **Critically endangered species** have a probability of extinction higher than 50% in ten years or three generations.
- **Endangered species** have a greater than 20% probability in 20 years or five generations.
- **Vulnerable species** have a probability greater than 10% over a century.

However, this is all based on 'business as usual' assumptions about how we will treat these species in the future. This assumption may not hold because classifying a species as 'endangered' or 'vulnerable' can be a call to action for conservation groups. For instance, a recent study of 48 bird and mammal populations found that many species classed as 'endangered' or 'vulnerable' had increasing populations, though many are still decreasing.

How many species has conservation saved?

An estimated 32 bird and 16 mammal species have been prevented from going extinct due to conservation efforts since 1993. They are shown by their extinction risk category, and status of how their populations were changing on the IUCN Red List in 2019.

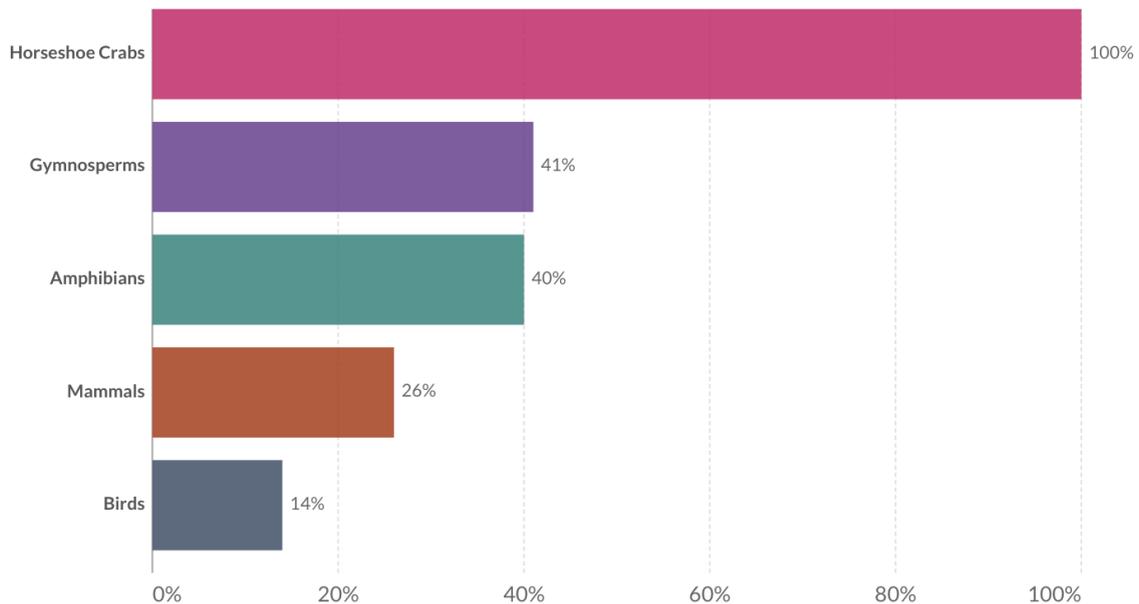


Source: Friederike Bolam et al. (2020). How many bird and mammal extinctions has recent conservation action prevented? *Conservation Letters*. OurWorldInData.org - Research and data to make progress against the world's largest problems. Licensed under CC-BY by the author Hannah Ritchie.

The following chart shows among the taxonomic groups for which at least 80% of species have been assessed, the share of species threatened with extinction.

Share of species threatened with extinction

Threatened species are those categorized as 'Critically Endangered', 'Endangered' and 'Vulnerable' on the IUCN Red List. This is shown by taxonomic group, and only for the more completely evaluated groups (where >80% of species have been evaluated).



Source: IUCN Red List

OurWorldInData.org/biodiversity • CC BY

CHART

TABLE

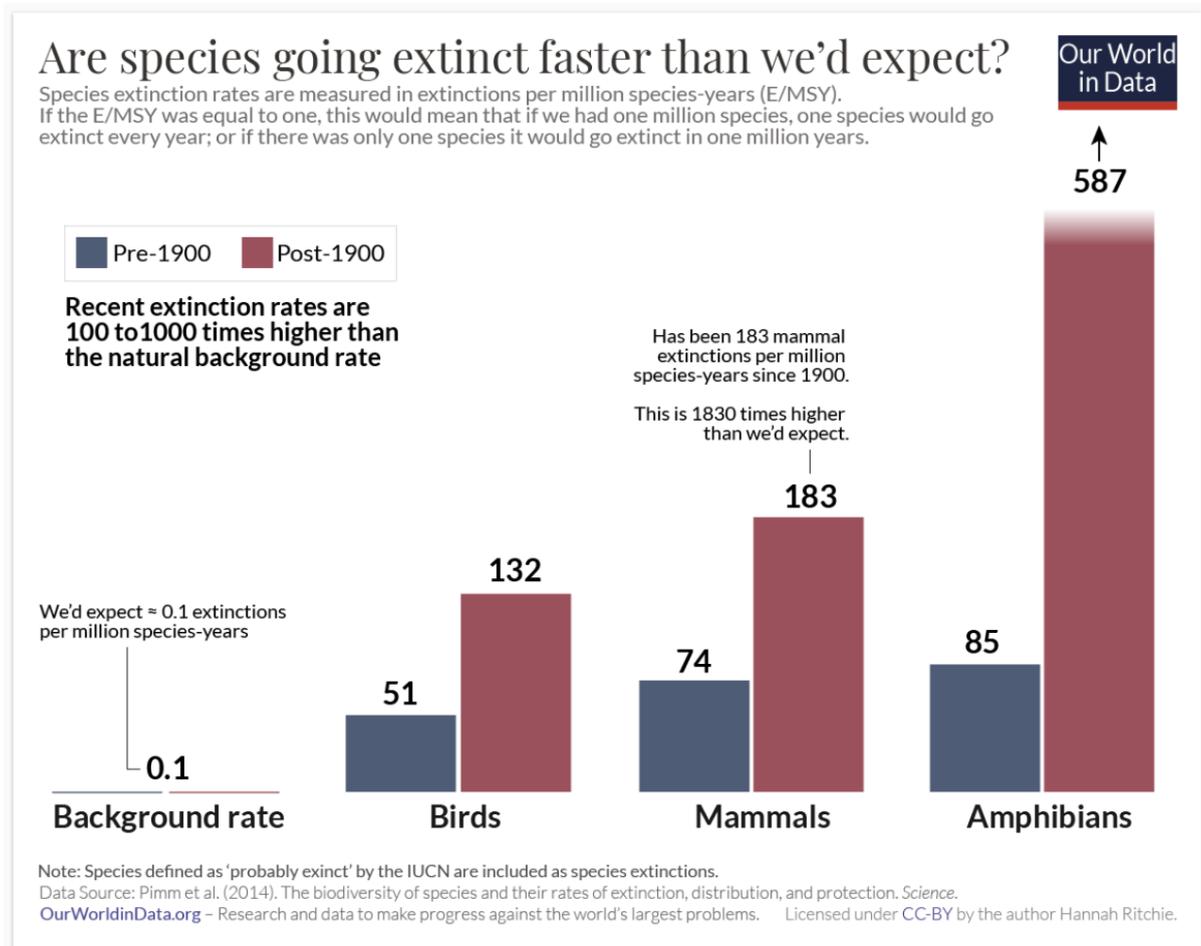
SOURCES

DOWNLOAD

Share

Due to elevated extinction rates, many scholars argue that we are in the middle of a Sixth Mass Extinction.

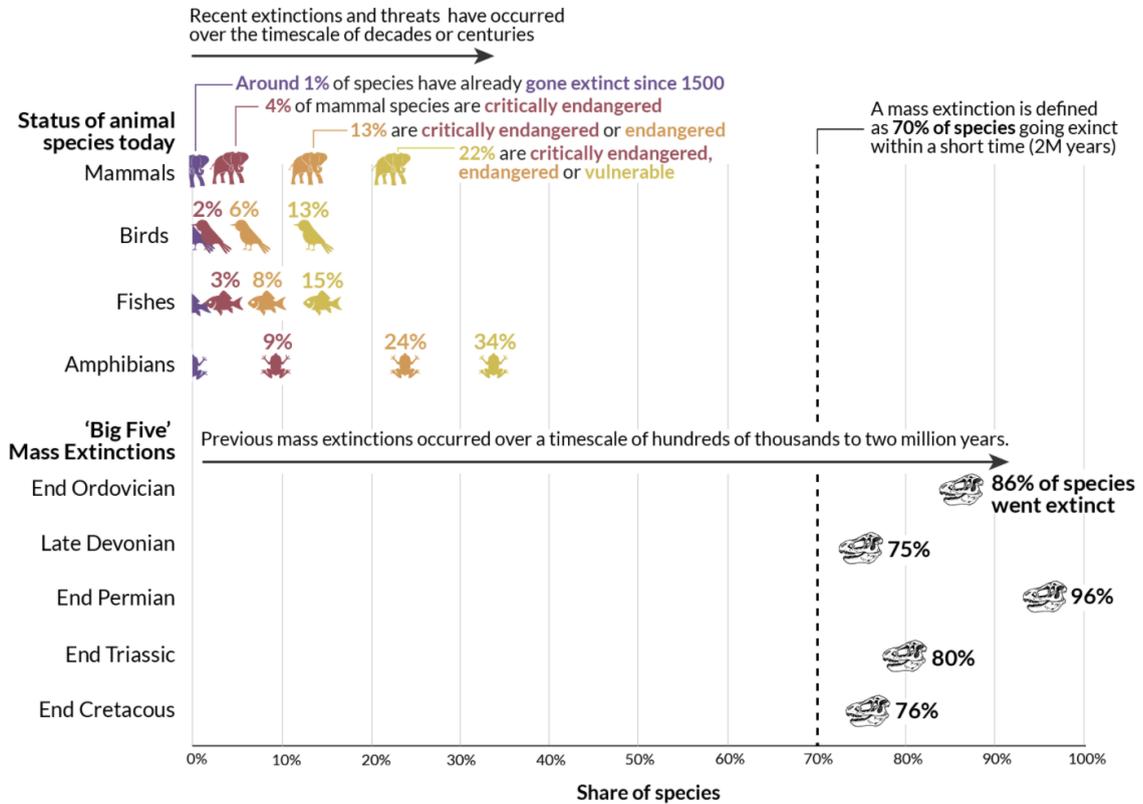
The rate of extinction is measured in extinctions per million-species years (E/MSY). The natural rate of extinction is 0.1 to 1 extinctions per million-species years. The rate of species extinctions is much faster than the natural background rate.



Since 1500, 1% of species have gone extinct, so we seem to be quite far from the 75% threshold. But we also need to consider the number of species *threatened* with extinction. Recall that vulnerable species have a probability of extinction greater than 10% over a century. If you assume that all vulnerable species will go extinct, then we are much closer to the 75% threshold.

How far are we from a sixth mass extinction?

Shown is the share of assessed animal species that have gone extinct or are threatened with extinction today, relative to the share of species that went extinct in previous mass extinction events. This is only shown for species in vertebrate groups where more than 80% of known species have been assessed for their extinction risk.



Data Sources: Barnosky et al. (2011). Has the Earth's sixth mass extinction already arrived? *Nature*. Threatened species from IUCN Red List (2021). Images sourced from Noun Project. OurWorldinData.org - Research and data to make progress against the world's largest problems. Licensed under CC-BY by the author Hannah Ritchie.

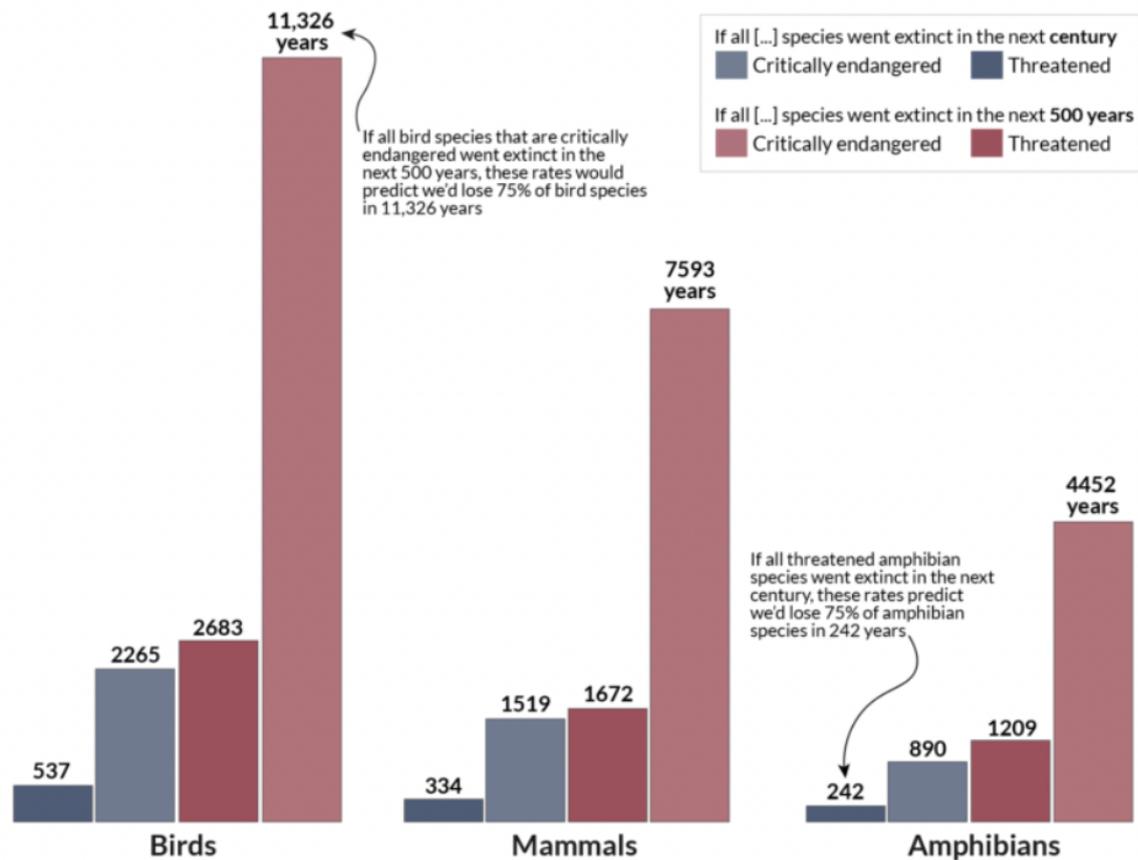
It is unclear how we should estimate the time it would take for there to be a sixth mass extinction. One option would be to base the estimate solely on the number of species recorded as extinct. Since 1500, we have lost 1% of species. If you project this average rate forward, then it would take 37,500 years to reach the 75% threshold. This is fast in geological terms but we would have lots of time to adapt and respond to this. If the post-1980 rate continues into the future, we would get there faster - in 18,000 years.

But this estimate would change if we make different assumptions about which species will go extinct:

How many years until we reach a sixth mass extinction?

Our World
in Data

A mass extinction is defined by a 75% loss of species within a short period of time (less than two million years). Years until we reach this point is shown under a number of hypothetical future scenarios with varying extinction rates.



Data Source: Anthony Barnosky et al. (2011). Has the Earth's sixth mass extinction already arrived? *Nature*.

OurWorldinData.org - Research and data to make progress against the world's largest problems.

Licensed under CC-BY by the author Hannah Ritchie.

The dark blue scenarios are pessimistic. Threatened species include 'vulnerable' species which have a 10% chance of going extinct within 100 years, so it would be a surprise if 100% went extinct over 100 years. The light red scenarios are optimistic. 'Critically endangered' species are defined as those that have a probability greater than 50% of going extinct in the next 50 years, so it would be a surprise if it took 500 years for 100% of them to go extinct, rather than 100 years.

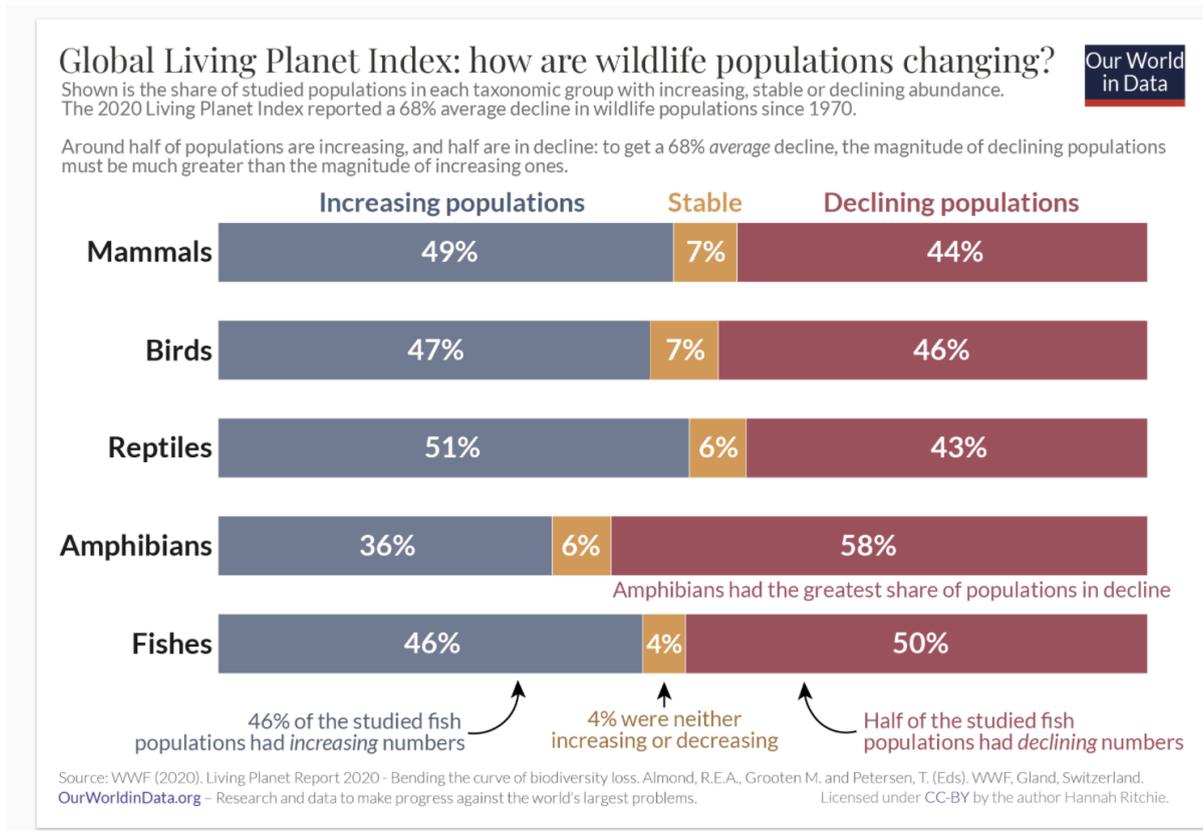
So, the light blue and dark red scenario seem more plausible: they suggest that it would take roughly 1,000 to 3,000 years for there to be a mass extinction, if current trends continue.

Population losses

Another argument that we are in a sixth mass extinction points to population losses.²³⁰ The Living Planet Report assesses trends in more than 20,000 wildlife populations, including 4,000 species across the world, which is only a small sample of the 2 million identified

²³⁰ Gerardo Ceballos, Paul R. Ehrlich, and Rodolfo Dirzo, 'Biological Annihilation via the Ongoing Sixth Mass Extinction Signaled by Vertebrate Population Losses and Declines', *Proceedings of the National Academy of Sciences* 114, no. 30 (2017): E6089–96.

species. The report is global in scope. The chart below shows which populations are increasing and which are decreasing.



The average decline across wildlife populations since 1970 is 68%. This means that the absolute magnitude of the decline is much greater in the declining populations than in the increasing populations. This doesn't mean that 68% of studied populations are declining. Our World in Data discusses the meaning of the average figure [here](#).

The import of this for the Sixth Mass Extinction is unclear. If species are evenly spread across the increasing and decreasing populations, then even if some populations go extinct, other populations would increase in abundance and so the whole species would not go extinct. But it might be that all of the species are concentrated on one side of the chart: maybe all rhino populations are declining, whereas all seagulls populations are increasing. The chart above does not tell us what is happening at the species level.

Speciation and hybridisation

Overall biodiversity is determined by the balance between extinction and speciation. Although extinctions have been increasing, speciation has also been increasing. Thomas (2015) explains why speciation has increased:

“The human-assisted movement of plants, animals, and microbes around the world has increased hugely over recent centuries, breaking down geographic barriers between species that exhibit incomplete genetic barriers to reproduction and, hence, setting the scene for a massive increase in levels of hybridisation [8–12]. Subsequent

genetic changes, including duplication of the entire genome (polyploidy) and chromosomal rearrangements, have, in one or a few generations, converted small numbers of these hybrid individuals or their offspring into sexually reproducing species that have limited compatibility with the parental species [9, 12–16]. The new hybrids can be at least as genetically distinct (by virtue of genomes derived from two parental species) as congeneric species that have arisen through geographic separation over longer periods of time.”²³¹

Direct genetic modification of plants is another way that humans can accelerate speciation.

Thomas (2015) estimates that current speciation could be 100 - 10,000 times faster than the natural background rate.²³² Data on both speciation and extinction is very poor, but according to Thomas (2019), as a rule, due to anthropogenic influence, local diversity stays about the same, regional diversity increases and global diversity declines.²³³ More new plant species have come into existence in Europe over the past three centuries than have been documented as becoming extinct over the same period, even though most new hybrid-origin species are likely to remain undetected.²³⁴ But at the global level species diversity has declined.

As I discuss below, ecosystem services are largely determined by the relationships within local ecosystems.

5.7.2. Drivers of species extinctions

The main posited drivers of species extinctions so far have been:²³⁵

- Land use change and habitat loss
- Direct exploitation or predation
- Sea-use change

Many scholars believe that climate change will have an increasingly large impact in the future.

²³¹ Chris D. Thomas, ‘Rapid Acceleration of Plant Speciation during the Anthropocene’, *Trends in Ecology & Evolution* 30, no. 8 (2015): 1. See also Bull, J. W., & Maron, M. (2016). How humans drive speciation as well as extinction. *Proceedings of the Royal Society B: Biological Sciences*, 283(1833), 20160600.

²³² “Considering these together, the Anthropocene plant speciation rate could be two to four orders of magnitude greater than the background rate.” Thomas, ‘Rapid Acceleration of Plant Speciation during the Anthropocene’, 6.

²³³ Chris D. Thomas, ‘The Development of Anthropocene Biotas’, *Philosophical Transactions of the Royal Society B: Biological Sciences* 375, no. 1794 (16 March 2020): sec. 4, <https://doi.org/10.1098/rstb.2019.0113>.

²³⁴ “Considering these together, the Anthropocene plant speciation rate could be two to four orders of magnitude greater than the background rate.” Thomas, ‘Rapid Acceleration of Plant Speciation during the Anthropocene’, 1.

²³⁵ Sandra Diaz et al., ‘The Global Assessment Report on Biodiversity and Ecosystem Services’, *The United Nations’ Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*, 2019, 28.

The Species-Area Relationship

Biodiversity loss is often estimated using the Species-Area Relationship, which is described as “ecology’s oldest law”.²³⁶ The Species-Area Relationship describes a widespread relationship between the area investigated and the number of species present. Bigger areas have more species, but the relationship between area and species number is curved: it typically follows a power function of the form $S = cA^z$, where S is the number of species, A is area, and c and z are constants. Thus, on log-transformed axes the relationship between S and A is linear and the slope of the best fitting regression line gives the exponent of the power function, z .

The charts below the species area relationship for a species of butterfly (note the log scale)

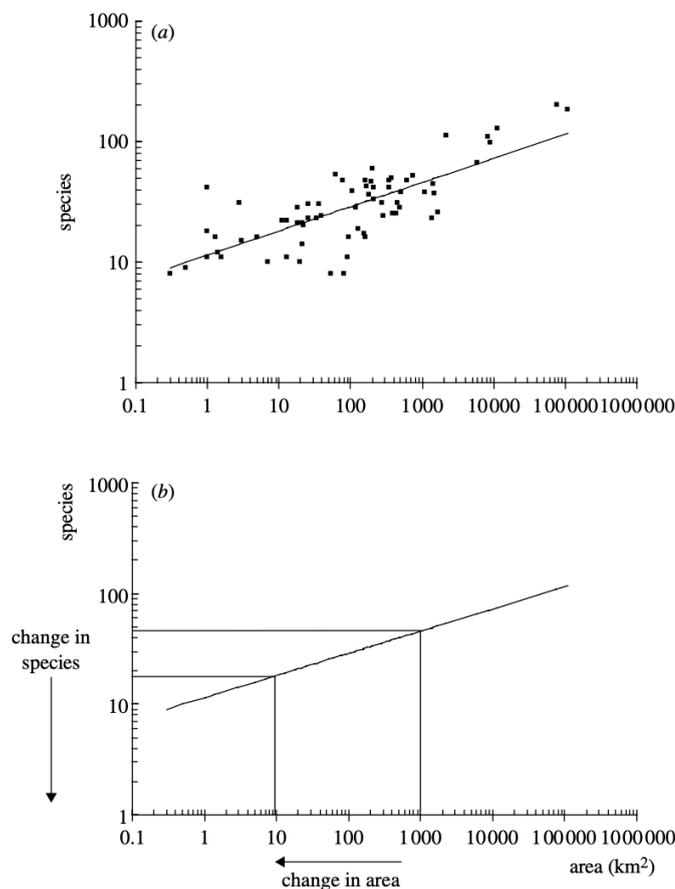


Figure 1. (a) Species–area relationship for butterflies on Caribbean Islands. Note the log-transformed axes. Redrawn from data in Davies & Spencer Smith (1997). (b) Estimating extinctions from habitat loss based on the species–area relationship. For a given reduction in area, the predicted loss of species depends on the slope of the line, z .

Source: Owen T Lewis, ‘Climate Change, Species–Area Curves and the Extinction Crisis’, *Philosophical Transactions of the Royal Society B: Biological Sciences* 361, no. 1465 (29 January 2006): 163–71, <https://doi.org/10.1098/rstb.2005.1712>

²³⁶ Owen T Lewis, ‘Climate Change, Species–Area Curves and the Extinction Crisis’, *Philosophical Transactions of the Royal Society B: Biological Sciences* 361, no. 1465 (29 January 2006): 163–71, <https://doi.org/10.1098/rstb.2005.1712>

Since bigger areas have more species, destroying habitat leading to some species going extinct or being 'committed to extinction', which means that species may survive for many generations.

Climate change as a driver

The Species Area Relationship has been applied to assess the potential impact of climate change. Since climate change alters the area and location of habitat available to a species, climate change will also cause extinctions, according to the Species Area Relationship.

One prominent pessimistic paper on the effect of climate change on ecosystems is by Thomas et al (2004). They assume that the 'climate envelope' available to species will move faster than species are able to disperse. How many species will go extinct as a result depends on species' ability to track shifting climates. Thomas et al (2004) estimate that due to climate change, 15-37% of all species will become 'committed to extinction' by mid-century.²³⁷ Since species are 'committed to' extinction and not actually extinct, there may be a lag of a few decades between a species being committed to extinction and actually going extinct.²³⁸

Most studies are less pessimistic than Thomas et al (2004). For example, the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services states:

“For instance, a synthesis of many studies estimates that the fraction of species at risk of extinction due to climate change is 5 per cent at 2°C warming, rising to 16 per cent at 4.3°C warming.”²³⁹

The percentages here are lower and 'risk of extinction' is a lower bar than 'committed to extinction', as the latter suggests that the extinction is guaranteed.

The IPCC provides the following diagram which illustrates how, according to some models, warming will be so fast on RCP8.5 that many species will not be able to move fast enough to stay in their ecological niche.

²³⁷ Chris D. Thomas et al., 'Extinction Risk from Climate Change', *Nature* 427, no. 6970 (January 2004): 145–48, <https://doi.org/10.1038/nature02121>.

²³⁸ “decades might elapse between area reduction (from habitat loss) and extinction” Chris D. Thomas et al., 'Extinction Risk from Climate Change', *Nature* 427, no. 6970 (January 2004): 145–48, <https://doi.org/10.1038/nature02121>.

²³⁹ “Globally, land-use change is the direct driver with the largest relative impact on terrestrial and freshwater ecosystems, while direct exploitation of fish and seafood has the largest relative impact in the oceans (well established) (Figure SPM.2) {2.2.6.2}. Climate change, pollution and invasive alien species have had a lower relative impact to date but are accelerating” Sandra Diaz et al., 'The Global Assessment Report on Biodiversity and Ecosystem Services', *The United Nations' Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*, 2019, 16.

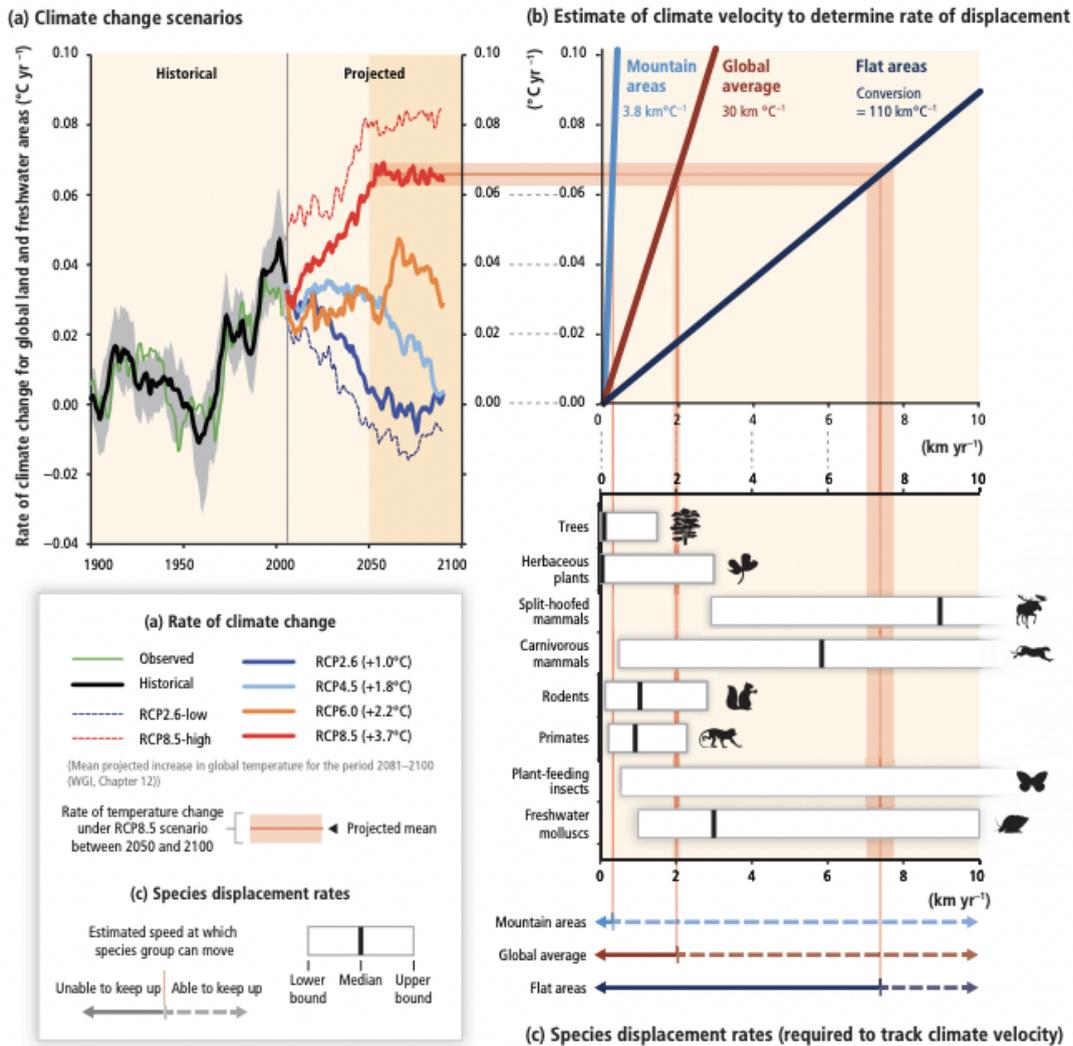


Figure 4-5 | (a) Rates of climate change, (b) corresponding climate velocities, and (c) rates of displacement of several terrestrial and freshwater species groups in the absence of human intervention. Horizontal and vertical pink bands illustrate the interpretation of this figure. Climate velocities for a given range of rates of climate change are determined by tracing a band from the range of rates in (a) to the points of intersection with the three climate velocity scalars in (b). Comparisons with species displacement rates are made by tracing vertical bands from the points of intersection on the climate velocity scalars down to the species displacement rates in (c). Species groups with displacement rates below the band are projected to be unable to track climate in the absence of human intervention. (a) Observed rates of climate change for global land areas are derived from Climatic Research Unit/Hadley Centre gridded land-surface air temperature version 4 (CRUTEM4) climate data reanalysis; all other rates are calculated based on the average of Coupled Model Intercomparison Project Phase 5 (CMIP5) climate model ensembles for the historical period (gray shading indicates model uncertainty) and for the future based on the four Representative Concentration Pathway (RCP) emissions scenarios. Data were smoothed using a 20-year sliding window, and rates are means of between 17 and 30 models using one member per model. Global average temperatures at the end of the 21st century for the four RCP scenarios are from WGI AR5 Chapter 12. (b) Estimates of climate velocity for temperature were synthesized from historical and projected future relationships between rates of temperature change and climate velocity (historical: Burrows et al., 2011; Chen et al., 2011; Dobrowski et al., 2013; projected future: Loarie et al., 2009; Sandel et al., 2011; Feeley and Rehm, 2012). The three scalars are climate velocities that are representative of mountainous areas (left), averaged across global land areas (center), and large flat regions (right). (c) Rates of displacement are given with an estimate of the median (black bars) and range (boxes = approximately 95% of observations or models for herbaceous plants, trees, and plant-feeding insects or median \pm 1.5 inter-quartile range for mammals). Displacement rates for herbaceous plants were derived from paleobotanical records, modern plant invasion rates, and genetic analyses (Kinlan and Gaines, 2003). Displacement estimates for trees are based on reconstructed rates of tree migration during the Holocene (Clark, 1998; Clark et al., 2003; Kinlan and Gaines, 2003; McLachlan et al., 2005; Nathan, 2006; Pearson, 2006) and modeled tree dispersal and establishment in response to future climate change (Higgins et al., 2003; Iverson et al., 2004; Epstein et al., 2007; Goetz et al., 2011; Nathan et al., 2011; Meier et al., 2012; Sato and Ise, 2012). Displacement rates for mammals were based on modeled dispersal rates of a wide range of mammal species (mean of Schloss et al., 2012 for Western Hemisphere mammals and rates calculated from global assessments of dispersal distance by Santini et al., 2013 and generation length by Pacifici et al., 2013). Displacement rates for phytophagous insects are based on observed dispersal distances and genetic analyses (Peterson and Denno, 1998; Kinlan and Gaines, 2003; Schneider, 2003; Berg et al., 2010; Chen et al., 2011). The estimate of median displacement rate for this group exceeds the highest rates on the axis. These displacement rates do not take into account limitations imposed by host plants. Displacement estimates for freshwater molluscs correspond to the range of passive plus active dispersal rates for upstream movement (Kappes and Haase, 2012).

Source: IPCC, *Climate Change 2014: Impacts, Adaptation, and Vulnerability* (Cambridge University Press, 2014), Fig. 4-5.

5.7.3. Could ecosystem collapse destroy agriculture?

The causal chain for the ecosystem collapse argument is as follows:

Global warming + habitat loss + human predation + pollution => global species loss
=> global ecosystem collapse => global agricultural catastrophe

I will now examine the plausibility of each part of this causal chain.

From human intervention to the destruction of agricultural ecosystem services

I will start by examining the first part of the causal chain

Global warming + habitat loss + human predation + pollution => global species loss

Before I discuss this part of the causal chain, it should be noted that the instrumental value of ecosystems is almost entirely a product of the relationships within local ecosystems. Global average loss of biodiversity matters only insofar as it is indicative of local biodiversity loss. Following the discovery of the Hawaiian Islands by the Polynesians 1500 years ago, they eliminated so many species that even the decadal *global* extinction rate would have been exceptional.²⁴⁰ This has next to no bearing on the risk of ecological collapse outside of Hawaii. (It also did not cause catastrophic ecosystem collapse *in* Hawaii).

Thus, damage to biodiversity only threatens global civilisation if local biodiversity is declining in all regions. I will now discuss several reasons that this appears not to be the case.

Extinctions have mainly occurred on islands

Of extinctions registered up to 2012, 95% occurred on islands or Australia and not on continents.²⁴¹ Since 1500, only six continental birds and three continental mammals have gone extinct. The rate of *continental* extinctions per million species years is therefore 1.61, which is below one estimate of the 'natural' or background rate of 2 extinctions per million species years (though this is likely an overestimate of the background rate).²⁴²

This is despite enormous deforestation. According to Loehle and Eschenbach (2012), "human predation (e.g. unregulated hunting and gathering of eggs) was a major factor in most confirmed extinctions, particularly on islands".

"However, it is worth noting that to date, no continental mammal or bird in our databases has been documented to have gone extinct solely because of habitat reduction. Early prehistoric waves of extinction (America around 12,000BP and Australia over 50,000 year ago) also were not because of habitat alteration (except fire) but largely because of hunting and other exploitation (e.g. egg gathering) (Trueman et al., 2005; Sodhi et al., 2009; Prideaux et al., 2010; Ripple & Van Valkenburgh, 2010). All extinctions of marine mammals are strictly because of

²⁴⁰ José M. Montoya, Ian Donohue, and Stuart L. Pimm, 'Planetary Boundaries for Biodiversity: Implausible Science, Pernicious Policies', *Trends in Ecology & Evolution* 33, no. 2 (1 February 2018): 71–73, <https://doi.org/10.1016/j.tree.2017.10.004>.

²⁴¹ Craig Loehle and Willis Eschenbach, 'Historical Bird and Terrestrial Mammal Extinction Rates and Causes', *Diversity and Distributions* 18, no. 1 (2012): 84–91, <https://doi.org/10.1111/j.1472-4642.2011.00856.x>.

²⁴² Gerardo Ceballos et al., 'Accelerated Modern Human-Induced Species Losses: Entering the Sixth Mass Extinction', *Science Advances* 1, no. 5 (1 June 2015): e1400253, <https://doi.org/10.1126/sciadv.1400253>.

hunting and are obviously not because of habitat alteration or introduced predators."²⁴³

This weakens one argument for the following causal links:

Habitat loss => global species loss

Climate change => global species loss

As we saw above, one argument maintains that because the rate of species extinctions is so high, we are facing a global mass extinction event if trends continue. But this does not follow because these extinctions are heavily concentrated among endemic species on islands. The rate of continental species extinctions is close to one estimate of the 'natural' background rate. The extinction of the dodo, which was endemic to Mauritius, is quite representative of extinctions in the last few centuries, but it is not relevant to ecosystem services outside of Mauritius.

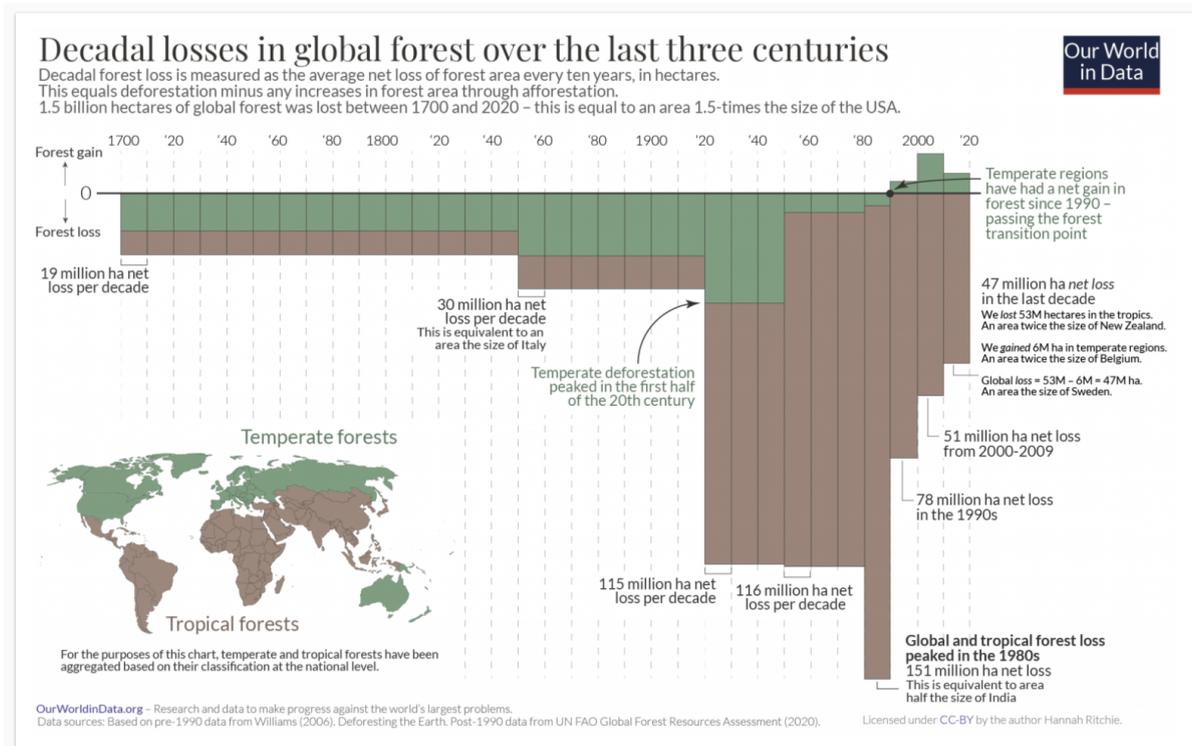
This also suggests that the species that are being killed off are not providing ecosystem services that are necessary for global agriculture. Animal extinctions are concentrated among endemic species on islands, rather than species that might be necessary to global agriculture.

Each of these points undermine one argument for the 'ecosystem collapse => global agricultural collapse' causal chain. However, other arguments for that causal chain might be more plausible.

Forest cover is increasing in many regions

Deforestation is declining in temperate regions. Net forest loss peaked in the 1980s and is declining.

²⁴³ Craig Loehle and Willis Eschenbach, 'Historical Bird and Terrestrial Mammal Extinction Rates and Causes', *Diversity and Distributions* 18, no. 1 (2012): 84–91, <https://doi.org/10.1111/j.1472-4642.2011.00856.x>.



Houghton and Nassikas (2017) estimate that forest area started increasing outside the tropics in 1950 not 1990, but that deforestation in the tropics peaked in 2000 rather than in the 1980s.²⁴⁴ I'm not sure what explains these discrepant findings.

There is also research suggesting that conventional estimates of global net forest change are overestimates because they neglect the countervailing effects of fire management, plantations and replanting. Accounting for this, Mendelsohn and Sohngen (2019) argue that net forest cover has increased substantially since 1900.²⁴⁵ This is just one study however and it is too soon to know whether opinion in the field will shift towards their position.

All of this is relevant to the following part of the causal chain:

Habitat loss => global species loss

Since forest habitat is increasing outside the tropics, habitat destruction is unlikely to be a driver of net species extinctions outside the tropics. If the extent of habitat loss is the main driver of biodiversity loss, then one would expect the biodiversity situation to improve in temperate regions. Thus, habitat loss does not threaten global agricultural catastrophe due to ecosystem collapse.

²⁴⁴ R. A. Houghton and Alexander A. Nassikas, 'Global and Regional Fluxes of Carbon from Land Use and Land Cover Change 1850–2015', *Global Biogeochemical Cycles* 31, no. 3 (2017): Fig. 1, <https://doi.org/10.1002/2016GB005546>.

²⁴⁵ Robert Mendelsohn and Brent Sohngen, 'The Net Carbon Emissions from Historic Land Use and Land Use Change', *Journal of Forest Economics* 34, no. 3–4 (2019): 263–83.

Local species richness is stable

Although it is declining globally, biodiversity is increasing or stable in many regions. This is because losses are being offset by increasing numbers of introduced or 'invasive' species. Various meta-analyses have shown that local species richness is fairly stable:

“Regardless of habitat type, geographic region, or most of the dominant ecological impacts discussed by the authors of individual studies (e.g., grazing, climate change, or pollution), the average change in species richness was not significantly different from zero (Vellend et al., 2013; Figure 4.2a). In a meta-analysis of this nature, one must always worry about biases in terms of where and when people have conducted empirical studies. For example, our sample of studies was distributed in a decidedly non-random way around the globe, with a preponderance of studies in Europe and North America—as is the case for essentially any general topic in ecology. However, in the 28 studies from under-represented parts of the globe (South America, Africa, Asia, and Australia), there was, if anything, a slight tendency for increases in species richness over time. To us, our results constituted a strong contradiction of the assumption that local-scale plant biodiversity loss is widespread and of large magnitude (e.g., >20%) in situations where biodiversity change can potentially impact ecosystem services.”

“Interestingly, our study was followed in quick succession by three other meta-analyses using independent data from different taxa and ecosystems showing essentially the same thing—no significant directional biodiversity change when averaged across many studies (Dornelas et al., 2014; Supp and Ernest, 2014)—or even an average increase in local-scale diversity in marine systems (Elahi et al., 2015; see Figure 4.2b,c). Taken together, to my eye, there is now clear evidence countering the notion of biodiversity declines at the local scale as a general, widespread rule.”²⁴⁶

The figure below shows there to be little change in species richness over time

²⁴⁶ Mark Vellend, 'Are Local Losses of Biodiversity Causing Degraded Ecosystem Function?', in *Effective Conservation Science*, by Peter Kareiva, Michelle Marvier, and Brian Silliman (Oxford University Press, 2017).

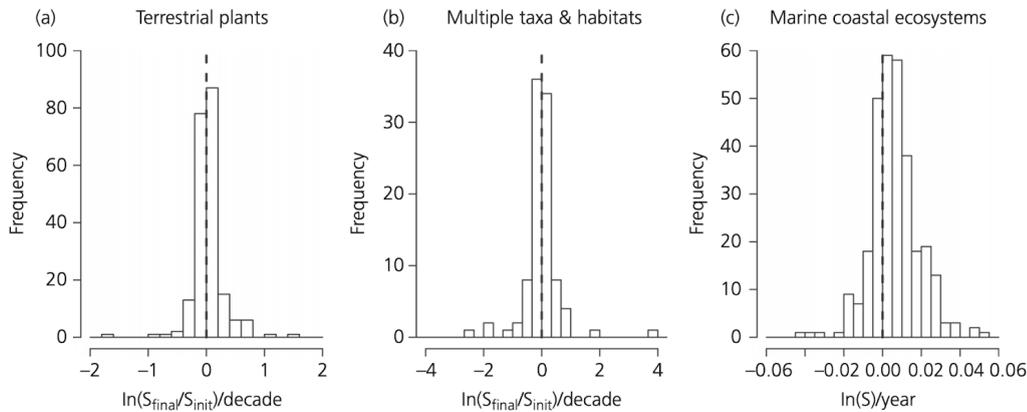


Figure 4.2 Distributions of temporal change in species richness (S) in three meta-analyses: (a) 212 data sets for terrestrial plants (Vellend et al., 2013; Vellend et al., 2017b); (b) 100 time series from a variety of terrestrial, freshwater, and marine ecosystems (Dornelas et al. 2014; two extreme positive values >5.5 not shown); (c) 302 time series in coastal marine ecosystems (Elahi et al., 2015). Temporal change is expressed as the log ratio of the final versus initial species richness in a and b and as the slope of species richness versus time in c. In all panels, the dashed line indicates zero temporal change. Data for a and b are publicly available in Vellend et al. (2017b) and for c in a github repository referenced in Elahi et al. (2015). Reprinted with permission from Vellend (2017).

Source: Mark Vellend, *Are Local Losses of Biodiversity Causing Degraded Ecosystem Function?*

One potential counter-argument to this is that although local biodiversity is stable, (1) the new local ecosystems have lower ecosystem services than the original ecosystems, and (2) the failure of these ecosystems would be more correlated, which poses a greater risk to global agriculture.

On (1), one thing we can say is that if lower ecosystem services have had an effect, it has not much effect so far given the massive increase in food yields over the last century. But it might be that there is some unknown nonlinear tipping point after which ecosystem services collapse. I discuss this possibility below. Independent of that, it is difficult to see why (1) would be true. A meta-analysis of the effect of ‘invasive’ species by Vilà et al (2011) found that “the magnitude and direction of the impact varied both within and between different types of impact... On average, abundance and diversity of the resident species decreased in invaded sites, whereas primary production and several ecosystem processes were enhanced.”²⁴⁷

Track record of predictions based on species area relationships

Many claims that habitat loss and climate change will cause biodiversity loss depend on the species area relationship. Some, though not all, of these predictions have a poor track record. The table below lists some of these predictions.

²⁴⁷ Montserrat Vilà et al., ‘Ecological Impacts of Invasive Alien Plants: A Meta-Analysis of Their Effects on Species, Communities and Ecosystems’, *Ecology Letters* 14, no. 7 (2011): 702–8.

Table 1 Estimated rates of extinction (after Reid 1992; Stork 1997, 1999)

Estimate	Taxon (<i>all taxa</i> unless otherwise stated)	% global loss per decade	Method of estimation	Reference
<i>(1) Species-area estimates</i>				
One million species between 1975 and 2000		4.00	Extrapolation of past exponentially increasing trend	Myers (1979)
15–20% of species between 1980 and 2000		8–11	Estimated species–area curve; forest loss based on Global 2000 projections	Lovejoy (1980)
50% loss of species by 2000 or soon after; 100% by 2010–2025		20–30	Various assumptions	Ehrlich and Ehrlich (1981)
9% extinction by 2000		7–8	Based on Lovejoy's calculations using updated forest loss estimates	Lugo (1988)
12% of plant species in neotropics, 15% of bird species in Amazon basin	Plants and birds	–	Species–area curve ($z = 0.25$)	Simberloff (1986)
2000 plant species loss per year in tropics and subtropics; 25% loss of species in 1985–2015	Plants	8–9	Loss of half the species in area likely to be deforested by 2015	Raven (1987, 1988)
At least 7% loss of plant species	Plants	7	Half of species lost over next decade in 10 'hot-spots' covering 3.5% of forest area	Myers (1988)
0.2–0.3% species loss per year		2–6	Half of rainforest species assumed lost in tropical rainforests to be local endemics and becoming extinct with forest loss	Wilson (1988, 1989)
2–13% loss between 1990 and 2015		1–5	Species–area curve ($0.15 < z < 0.35$); range includes current rate of forest loss and 50% increase	Reid (1992)
37–50% extinction rate for 5,308 Amazonian plant species by 2020	Amazonian plant species	14.6–19.7	Estimated number of Amazonian species lost by Laurance et al. (2001) predicted forest loss	Hubbell et al. (2008)
5–9% extinction for Amazonian plants by 2050	Amazonian plant species	1.2–2.2	Revision of above estimate reviewing potential habitat loss and species distributions	Feeley and Silman (2009)
Predict: in next 25 years rate of net tropical deforestation to slow on all continents; net increase in forest area in Latin America and Asia at least within 50 years, and in Africa within 100 years			Slowing shifts of people from tropical rural to urban; suggested reduced deforestation, increased secondary forests succession, both reducing extinction pressures	Wright and Muller-Landau (2006a, b)

Table 1 continued

Estimate	Taxon (<i>all taxa</i> unless otherwise stated)	% global loss per decade	Method of estimation	Reference
<i>(2) Estimates from empirical data</i>				
Past loss of 50–90% Pacific Island birds	Pacific island birds		Evidence from recent fossil evidence	Pimm et al. (1995), Steadman (1995)
Current extinction rates 100 times background rates and will increase 15 fold by 2100	All bird species		Range of threats to bird species: habitat/forest loss, invasive species and global change	Pimm et al. (2006)
<i>(3) Estimates from Red lists</i>				
50% extinction in 50–100 years (palms and 300–400 years (birds and mammals)	Selected taxa including palms, birds and mammals	1–10	Extrapolating current recorded extinction rates and by dynamics of threatened status	Smith et al. (1993)
6–50% extinction of selected vertebrates in next 100 years	Some vertebrate taxa including reptile, mammal and bird groups	0.6–5	Fitting of exponential extinction functions based on IUCN categories of threat	Mace (1994)
100,000–500,000 insect species extinctions in next 300 years	Insects	0.04–0.21	Extrapolating from British Red Data Books and assuming 8 million species of insects worldwide	Mawdsley and Stork (1995)
<i>(4) Co-extinctions</i>				
3–5% extinction of helminth species parasitising vertebrates in next 50–100 years	Parasitic helminths	0.15–0.5	Loss of parasites due to loss of hosts; developing co-extinction theory (Stork and Lyal 1993; Koh et al. 2004)	Dobson et al. (2008)
213,830–547,500 insect herbivores committed to extinction in 34 biodiversity hot-spots	Herbivorous insects		Assumes that 150,371 endemic plant species have an average of 5.3–10.6 monophagous insect herbivores	Fonseca (2009)
<i>(5) Models of climate change</i>				
15–37% loss of species by 2050	1,103 species including vertebrates & invertebrates	3.2–7.8	Modelled species distributions against IPCC climate change scenarios	Thomas et al. (2004)

Note that many of these estimated rates include species 'committed' to extinction. Estimates of % global loss per decade are from Reid (1992) or for later estimates are calculated by the current author based on global estimates for all species on Earth of 10 million and for all insects of 8 million

Source: Nigel E. Stork, 'Re-Assessing Current Extinction Rates', *Biodiversity and Conservation* 19, no. 2 (1 February 2010): 357–71, <https://doi.org/10.1007/s10531-009-9761-9>.

With the benefit of hindsight, all of the estimates from the 1980s are substantial overestimates. If these estimates were correct, we would expect 25% to 100% of species to be committed to extinction today, which is far higher than estimates of 'critically endangered' species today.

The table also shows the estimate from Thomas et al (2004), which entails an implied species loss of 3.2% to 7.8% per decade. Since Thomas et al (2004) was published 17 years ago, his prediction would suggest that due to climate change, 7.5% of species are already committed to extinction due to climate change. This is at odds with documented extinctions so far. As I have mentioned, of the confirmed extinctions, 95% have occurred on islands due to direct predation, not due to global climate change.

Predictions of biodiversity loss from the species-area relationship have been the subject of criticism. Botkin et al (2007) outline six limitations of the species-area relationship as a tool to model biodiversity loss:

“First, it assumes an equilibrium (or very slowly changing) relationship between species number and area.

Second, the future climate probably will not be an exact analog of the current one, so “moving” a bioclimatic zone for an ecological type may not be accurate (Malcolm et al. 2006).

Third, topographic variation, which affects the species–area curve shape, may be greater or less in the future zone.

Fourth, factors relating to the shape of areas and the amount of their fragmentation suggest that an alternative “endemics–area curve” may enable more accurate predictions (Harte et al. 2004).

Fifth, the correct z value must be chosen: It must apply to the entire area under consideration, and it must also consider the type of area and timescale applicable (Rosenzweig 1995).

Sixth, many species are not confined to a particular vegetation zone or type. For the species–area relationship to predict species extinctions, the area must be for closed communities. Thomas and colleagues (2004) used individual species distributions as the basis for their analysis. They examined changes in realized niches without taking into account the likelihood of changed interactions and adaptation, and thus the new areas that they predicted were probably too small. How these area changes relate to changes in area of closed communities is unclear.”²⁴⁸

He and Hubbell (2011) argue that species–area relationships always overestimate extinction rates from habitat loss:

²⁴⁸ Daniel B. Botkin et al., ‘Forecasting the Effects of Global Warming on Biodiversity’, *BioScience* 57, no. 3 (1 March 2007): 227–36, <https://doi.org/10.1641/B570306>. See also Carsten F. Dormann, ‘Promising the Future? Global Change Projections of Species Distributions’, *Basic and Applied Ecology* 8, no. 5 (3 September 2007): 387–97, <https://doi.org/10.1016/j.baae.2006.11.001>; Owen T Lewis, ‘Climate Change, Species–Area Curves and the Extinction Crisis’, *Philosophical Transactions of the Royal Society B: Biological Sciences* 361, no. 1465 (29 January 2006): 163–71, <https://doi.org/10.1098/rstb.2005.1712>.

“The key mathematical result is that the area required to remove the last individual of a species (extinction) is larger, almost always much larger, than the sample area needed to encounter the first individual of a species, irrespective of species distribution and spatial scale. We illustrate these results with data from a global network of large, mapped forest plots and ranges of passerine bird species in the continental USA; and we show that overestimation can be greater than 160%”²⁴⁹

Models vs the paleoclimate

In section 3, I discussed at length how well ecosystems fared during periods in which temperatures were much higher than today, and warming was comparable fast on a regional basis, compared to today. Since the breakup of Pangea, climate change has not been correlated with elevated rates of species extinctions.

In contrast to this, the recent IPBES report predicts substantial species extinctions: 5% of species at risk of extinction at 2°C, rising to 16% at 4°C.²⁵⁰ As we saw above, this was on the basis of the Species-Area Relationship, which predicts that species’ ecological niche will move faster than their ability to disperse. This same model therefore predicts that there would have been extensive species loss during the transition from the Pleistocene to the Holocene.

As we saw in Chapter 3, that is not what happened. Despite regional warming of 2°C to 15°C per century, there is little evidence of local extinction (aka ‘extirpation’).²⁵¹ This strongly suggests that the models are wrong.

Willis and MacDonald (2011) list several reasons that might explain where the models go wrong.

1. **CO₂ fertilisation** - Higher levels of CO₂ increase photosynthesis and carbon uptake, “which may account for increased growth rates in the African and American rainforests in the last 30 years.” “Recent modeling has demonstrated that increased primary productivity has a linear positive relationship with diversity capacity such that a given ecosystem can support a greater number of species during intervals of higher CO₂ (Woodward & Kelly 2008). Interestingly, when results from this model are

²⁴⁹ Fangliang He and Stephen P. Hubbell, ‘Species–Area Relationships Always Overestimate Extinction Rates from Habitat Loss’, *Nature* 473, no. 7347 (2011): 368–71.

²⁵⁰ “Globally, land-use change is the direct driver with the largest relative impact on terrestrial and freshwater ecosystems, while direct exploitation of fish and seafood has the largest relative impact in the oceans (well established) (Figure SPM.2) {2.2.6.2}. Climate change, pollution and invasive alien species have had a lower relative impact to date but are accelerating” Sandra Diaz et al., ‘The Global Assessment Report on Biodiversity and Ecosystem Services’, *The United Nations’ Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*, 2019, 16.

²⁵¹ Terence P. Dawson et al., “Beyond Predictions: Biodiversity Conservation in a Changing Climate,” *Science* 332, no. 6025 (April 1, 2011): 53–58, <https://doi.org/10.1126/science.1200303>; Christian Hof et al., “Rethinking Species’ Ability to Cope with Rapid Climate Change,” *Global Change Biology* 17, no. 9 (September 1, 2011): 2987–90, <https://doi.org/10.1111/j.1365-2486.2011.02418.x>; K. J. Willis and G. M. MacDonald, “Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World,” *Annual Review of Ecology, Evolution, and Systematics* 42, no. 1 (2011): 267–87, <https://doi.org/10.1146/annurev-ecolsys-102209-144704>; Daniel B. Botkin et al., “Forecasting the Effects of Global Warming on Biodiversity,” *BioScience* 57, no. 3 (March 1, 2007): 227–36, <https://doi.org/10.1641/B570306>.

applied to future climatic scenarios, the output predicts enhanced plant growth, an increase in ecosystem productivity, and higher diversity (Woodward 2010). What is also apparent from this model, however, is that “weed” species for which migration is an insignificant barrier are likely to fill future diversity capacities”.

2. **Higher levels of ecological tolerance** - Many species had a much wider ecological tolerance than is apparent from their present day distributions, and thus they contain gene variations that enable tolerance of much higher temperatures and water stress.
3. **Plants survived in refugia** - “those plants unable to adapt became restricted to small, microenvironmentally favorable refugia where they were able to persist”.
“Another factor leading to persistence appears to have been the survival of populations in small, environmentally favorable refugial localities, as seen for many European alpine species during the mid-Holocene climatic optimum”

A common response to this is to argue that species will today have to adapt in the context of fragmented habitat, which gives less scope for dispersal. For instance, the IPCC says:

“Finally, evidence from the paleontological record indicating very low extinction rates over the last several hundred thousand years of substantial natural fluctuations in climate—with a few notable exceptions such as large land animal extinctions during the Holocene—has led to concern that forecasts of very high extinction rates due entirely to climate change may be overestimated (Botkin et al., 2007; Dawson et al., 2011; Hof et al., 2011a; Willis and MacDonald, 2011; Moritz and Agudo, 2013). However, as indicated in Section 4.2.3, no past climate changes are precise analogs of future climate change in terms of speed, magnitude, and spatial scale; nor did they occur alongside the habitat modification, overexploitation, pollution, and invasive species that are characteristic of the 21st century. Therefore the paleontological record cannot easily be used to assess future extinction risk due to climate change”²⁵²

It may be true that climate change is happening in the context of habitat modification, overexploitation, pollution and invasive species, but the models predicting substantial species loss due to climate *do not* factor this in. Therefore, the paleoclimatic evidence does indeed show that the models predicting substantial species loss are not reliable.

This does *not* necessarily mean that there will *not* be substantial species loss due to climate change. It is just to point out that the current models are mistaken. Indeed, it is plausible that because it will occur in the context of habitat modification, climate change will cause biodiversity loss.

Global species loss and agriculture

So far, I have focused on the first part of the causal chain

Global warming + habitat loss + human predation + pollution => global species loss

I am now going to focus on the later parts of the causal chain

²⁵² IPCC, *Climate Change 2014: Impacts*, 301.

Global species loss => global ecosystem collapse => global agricultural catastrophe

Clarifying the argument

It is worth clarifying exactly what is posited to happen in this part of the causal chain. The idea is that there would be a loss of species which are crucial to the production of all major food crops in all major food producing regions. This would be a very dramatic event, and I have not seen a detailed description of what exactly is meant to happen.

One example would be if pollinators like bees and butterflies went extinct due to climate change. This is discussed by Our World in Data [here](#). A third of crop production depends on pollinators; staple cereals that account for the majority of food production do not depend on pollinators. Moreover, few crops are entirely dependent on pollinators. According to Our World in Data, “studies suggest crop production would decline by around 5% in higher income countries, and 8% at low-to-middle incomes if pollinator insects vanished”.

Outside of crops dependent on pollinators, it is difficult to see how species loss could threaten destruction of all major crops in all major food producing regions in very diverse ecosystems and climates.

Closeness to pre-modern ecology and living standards are negatively correlated

Many scholars writing on biodiversity loss seem to believe that any departures from the pre-agricultural or pre-industrial ecosystem are bad. The Biodiversity Intactness Index measures such departures and is “defined as the average abundance of a taxonomically and ecologically broad set of species in an area, relative to their abundances in an intact reference ecosystem”, where the reference is usually the pre-industrial ecosystem.²⁵³ However, closeness to pre-industrial ecology appears to be negatively correlated with levels of consumption now and into the long-term.

Steffen et al (2015) argued that the planetary boundary for the Biodiversity Intactness Index is 90%.²⁵⁴ What exactly the implications of passing this alleged planetary boundary are meant to be is somewhat unclear, but one natural interpretation is that we risk passing a threshold which would diminish consumption substantially or threaten human civilisation. Indeed, the IPBES interprets the threshold in this way:

“That framework suggests that large regions whose biotic integrity – i.e., the fraction of originally-present biodiversity that remains – falls below 90% risk large-scale failure of ecosystem resilience that would cause critical reductions in the flows of nature’s contributions to people (Steffen et al. 2015b) though there is a great deal of uncertainty about precisely where any boundary should be placed”²⁵⁵

²⁵³ Samantha L. L. Hill et al., ‘Worldwide Impacts of Past and Projected Future Land-Use Change on Local Species Richness and the Biodiversity Intactness Index’, *BioRxiv*, 1 May 2018, 311787, <https://doi.org/10.1101/311787>.

²⁵⁴ Will Steffen et al., ‘Planetary Boundaries: Guiding Human Development on a Changing Planet’, *Science* 347, no. 6223 (13 February 2015): 1259855, <https://doi.org/10.1126/science.1259855>.

²⁵⁵ Sandra Diaz et al., ‘The Global Assessment Report on Biodiversity and Ecosystem Services’, *The United Nations’ Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*, 2019, chap. 2.2.

The 90% boundary figure seems arbitrary. Steffen et al give no reason that it should be 90%, rather than 1% or 99% or that there is no meaningful boundary.

“Due to a lack of evidence on the relationship between BII and Earth-system responses, we propose a preliminary boundary at 90% of the BII but with a very large uncertainty range (90 to 30%) that reflects the large gaps in our knowledge about the BII–Earth-system functioning relationship.”

I think proposing that there is a planetary boundary of 90% with no argument or evidence is unjustifiable. Brook et al (2013) argue that there is reason to think that there is no global level planetary boundary for biodiversity.

“By evaluating potential mechanisms and drivers, we conclude that spatial heterogeneity in drivers and responses, and lack of strong continental interconnectivity, probably induce relatively smooth changes at the global scale, without an expectation of marked tipping patterns. This implies that identifying critical points along global continua of drivers might be unfeasible and that characterizing global biotic change with single aggregates is inapt.”²⁵⁶

For further criticism of the idea of planetary boundaries see Nordhaus et al (2012).²⁵⁷

Many regions today have a Biodiversity Intactness Index of less than 90%.

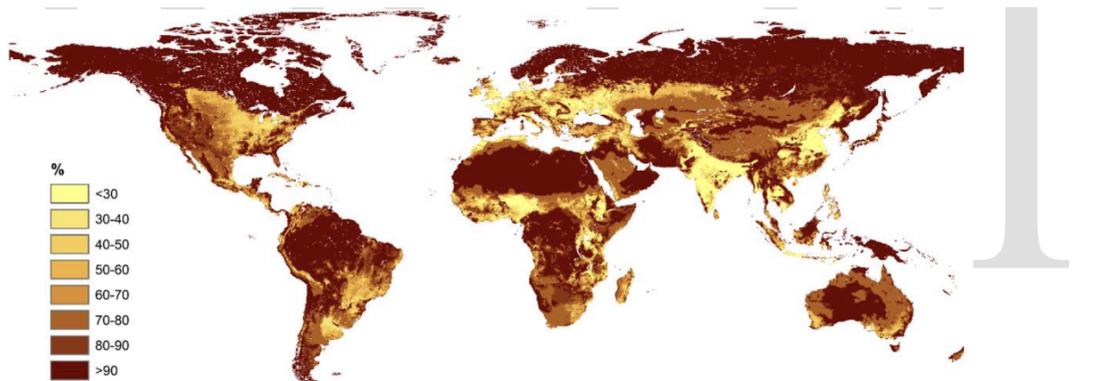


Figure 2.14_COMPOSITION: Global map of estimated terrestrial Biodiversity Intactness Index in 2015 (Hill et al. 2018). Darker colours indicate more intact ecological community composition.

Source: Sandra Diaz et al., ‘The Global Assessment Report on Biodiversity and Ecosystem Services’, The United Nations’ *Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*, 2019.

It is notable that many of the regions with the lightest colours are some of the richest in the world, including the US and many European countries. Many of the poorest countries in the

²⁵⁶ Barry W. Brook et al., ‘Does the Terrestrial Biosphere Have Planetary Tipping Points?’, *Trends in Ecology & Evolution* 28, no. 7 (1 July 2013): 396–401, <https://doi.org/10.1016/j.tree.2013.01.016>.

²⁵⁷ Ted Nordhaus, Michael Shellenberger, and Linus Blomqvist, ‘The Planetary Boundaries Hypothesis: A Review of the Evidence’ (Breakthrough Institute, June 2012), https://static1.squarespace.com/static/5c9e9320348cd97cfd4c123c/t/5e1a7a7a3a904b6e1b424be0/1578793595305/Planetary_Boundaries.pdf.

world score well on the Biodiversity Intactness Index, including for example DRC, Angola, and the Central African Republic.

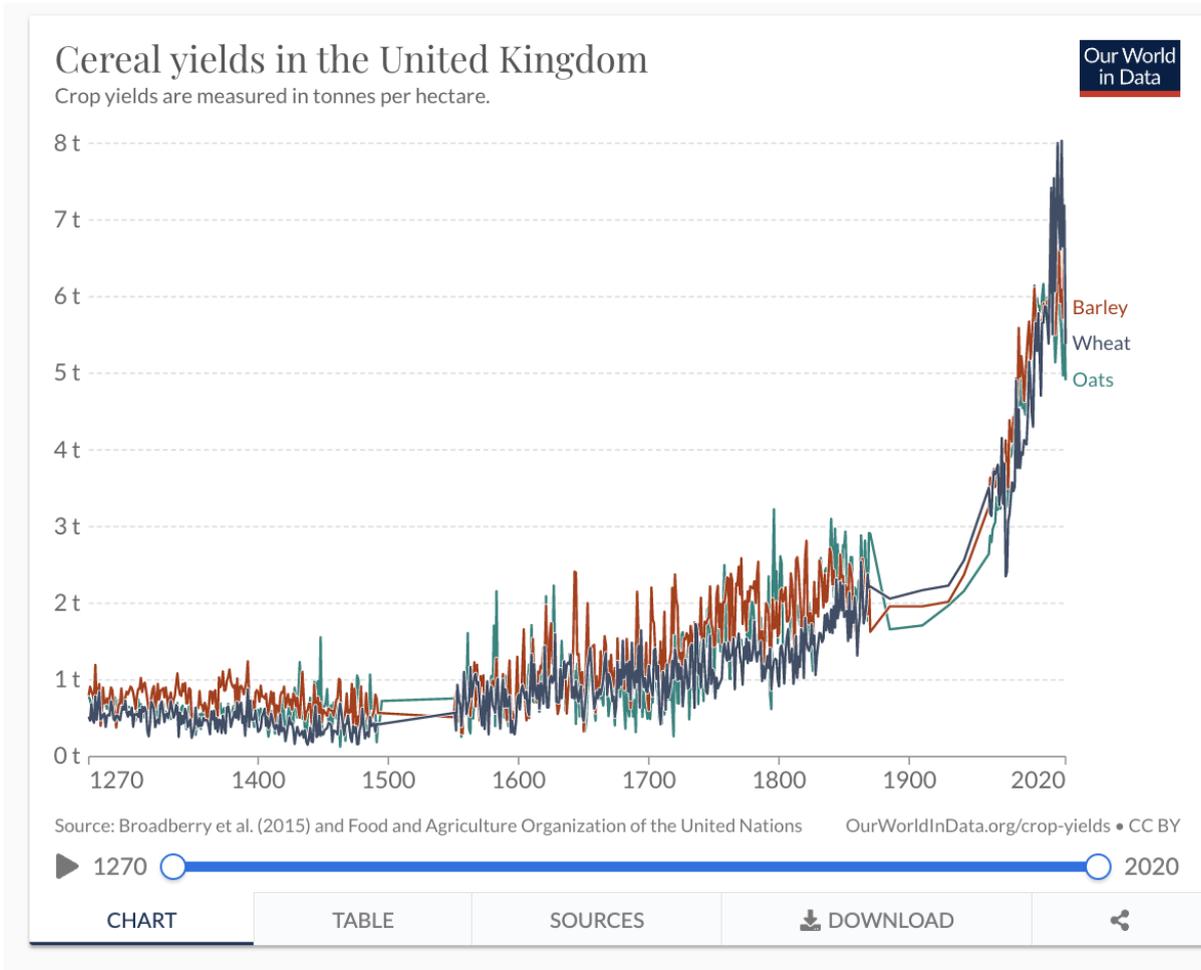
Clearly then, the claim cannot be that declines in Biodiversity Intactness instantaneously bring about declines in welfare. Rather, the claim must be that the decline in Biodiversity Intactness brings about a decline in welfare at some point in the future that overwhelms any gains in welfare prior to that, perhaps due to passing some unknown tipping point.

But this seems to be the opposite of the truth for many regions. If one were to bet which regions would have the highest levels of consumption in 300 years, one would probably say the US, Europe and China, all areas where the biodiversity intactness index is well below 90%. Indeed, extremely rich countries such as England, France and Denmark seem to have Index scores below even the Steffen et al 30% lower planetary boundary uncertainty range. The Biodiversity Intactness Index just does not seem to be a good measure of the long-term consumption prospects of different regions.

There is also strong evidence from history that complete destruction of ecosystems is very unlikely to trigger tipping points that do severe damage to society. Significant deforestation started across Eurasia thousands of years ago. According to one estimate, in England in 1,000 BC, 90% of land potentially suited to agriculture was covered with forest. By 1400 AD, this had plummeted to 17%. By 1850 AD, forest cover had fallen further to only 2%.²⁵⁸

At no point throughout this period did England experience nonlinear ecological collapse. Agricultural production in England increased enormously despite the complete destruction of pre-industrial ecosystems.

²⁵⁸ Jed O. Kaplan, Kristen M. Krumhardt, and Niklaus Zimmermann, 'The Prehistoric and Preindustrial Deforestation of Europe', *Quaternary Science Reviews* 28, no. 27–28 (2009): Table 3.



There is a similar picture across the rest of Europe.²⁵⁹ While pre-modern and pre-industrial ecosystems have been destroyed, living standards have increased enormously.

Raudsepp-Hearne et al (2010) call this the 'environmentalist's paradox: human well-being has increased despite large global declines in most ecosystem services.'²⁶⁰

Overall judgement on ecosystem collapse and agriculture

The causal chain for the ecosystem collapse argument is as follows:

Global warming + habitat loss + human predation + pollution => global species loss
=> global ecosystem collapse => global agricultural catastrophe

Each part of this causal chain seems to be flawed in several important ways. In my view, the risks of ecosystem collapse to global agricultural production are minimal.²⁶¹ This does not

²⁵⁹ Jed O. Kaplan, Kristen M. Krumhardt, and Niklaus Zimmermann, 'The Prehistoric and Preindustrial Deforestation of Europe', Table 3.

²⁶⁰ Ciara Raudsepp-Hearne et al., 'Untangling the Environmentalist's Paradox: Why Is Human Well-Being Increasing as Ecosystem Services Degrade?', *BioScience* 60, no. 8 (1 September 2010): 576–89, <https://doi.org/10.1525/bio.2010.60.8.4>.

²⁶¹ Kareiva and Carranza reach a similar conclusion: "The interesting question is whether any of the planetary thresholds other than CO2 could also portend existential risks. Here the answer is not clear. One boundary often mentioned as a concern for the fate of global civilization is biodiversity (Ehrlich &

mean that global ecosystem damage does not matter. My focus here has been on the instrumental value of ecosystems to human society.

Ehrlich, 2012), with the proposed safety threshold being a loss of greater than .001% per year (Rockström et al., 2009). There is little evidence that this particular .001% annual loss is a threshold—and it is hard to imagine any data that would allow one to identify where the threshold was (Brook et al., 2013; Lenton & Williams, 2013). A better question is whether one can imagine any scenario by which the loss of too many species leads to the collapse of societies and environmental disasters, even though one cannot know the absolute number of extinctions that would be required to create this dystopia. While there are data that relate local reductions in species richness to altered ecosystem function, these results do not point to substantial existential risks. The data are small-scale experiments in which plant productivity, or nutrient retention is reduced as species number declines locally (Vellend, 2017), or are local observations of increased variability in fisheries yield when stock diversity is lost (Schindler et al., 2010). Those are not existential risks. To make the link even more tenuous, there is little evidence that biodiversity is even declining at local scales (Vellend et al 2017; Vellend et al., 2013). Total planetary biodiversity may be in decline, but local and regional biodiversity is often staying the same because species from elsewhere replace local losses, albeit homogenizing the world in the process. Although the majority of conservation scientists are likely to flinch at this conclusion, there is growing skepticism regarding the strength of evidence linking trends in biodiversity loss to an existential risk for humans (Maier, 2012; Vellend, 2014). Obviously if all biodiversity disappeared civilization would end—but no one is forecasting the loss of all species. It seems plausible that the loss of 90% of the world's species could also be apocalyptic, but not one is predicting that degree of biodiversity loss either. Tragic, but plausible is the possibility our planet suffering a loss of as many as half of its species. If global biodiversity were halved, but at the same time locally the number of species stayed relatively stable, what would be the mechanism for an end-of-civilization or even end of human prosperity scenario? Extinctions and biodiversity loss are ethical and spiritual losses, but perhaps not an existential risk.” Peter Kareiva and Valerie Carranza, ‘Existential Risk Due to Ecosystem Collapse: Nature Strikes Back’, *Futures*, 5 January 2018, <https://doi.org/10.1016/j.futures.2018.01.001>.

6. Heat stress

Heat stress is dependent on temperature, humidity, wind speed and metabolic heat generation. Even though the hottest temperatures occur in subtropical deserts, relative humidity there is so low that maximal annual heat stress is no higher than in the deep tropics.²⁶² In more humid climates, sweating is less effective at reducing our body temperature.

There are two main ways that increasing heat stress could affect human society: by reducing labour capacity and increasing heat-related morbidity and mortality.

6.1. Metrics

There are numerous heat stress metrics used at the moment, all of which combine temperature and humidity. The most common in climate change research are Wet Bulb Temperature and Wet Bulb *Globe* Temperature.

Wet Bulb Temperature = The temperature that an air parcel would reach through evaporative cooling once fully saturated.²⁶³ Or, in layman's terms, the reading from a thermometer when covered in a wet cloth and swung in the air.

Wet Bulb *Globe* Temperature = A function of Wet Bulb Temperature, dry bulb temperature (what you would see on an ordinary thermometer), and a black globe thermometer (that measures the effect of solar radiation). It is $(0.7 \times \text{Wet Bulb} + 0.1 \times \text{Dry Bulb} + 0.2 \times \text{Black Globe Temperature})$.

Wet Bulb Temperature is usually used to measure extreme survivability limits, whereas Wet Bulb **Globe** Temperature is usually used to measure safe limits on activities for exposed people. (I am going to put 'globe' in bold to avoid confusion).

The table below shows Wet Bulb **Globe** Temperature limits on activities for acclimatised people.²⁶⁴ The limit is about 2°C lower for non-acclimatised people.²⁶⁵

Level of activity	WBGT limit for acclimatised people
Resting	34
Walking/hammering	28

²⁶² Steven C. Sherwood and Matthew Huber, 'An Adaptability Limit to Climate Change Due to Heat Stress', *Proceedings of the National Academy of Sciences* 107, no. 21 (25 May 2010): 9552–55, <https://doi.org/10.1073/pnas.0913352107>.

²⁶³ Ethan D. Coffel, Radley M. Horton, and Alex de Sherbinin, 'Temperature and Humidity Based Projections of a Rapid Rise in Global Heat Stress Exposure during the 21st Century', *Environmental Research Letters* 13, no. 1 (December 2017): 014001, <https://doi.org/10.1088/1748-9326/aaa00e>.

²⁶⁴ David Newth and Don Gunasekera, 'Projected Changes in Wet-Bulb Globe Temperature under Alternative Climate Scenarios', *Atmosphere* 9, no. 5 (May 2018): fig. 5, <https://doi.org/10.3390/atmos9050187>.

²⁶⁵ Newth and Gunasekera, fig. 5.

Intense - running, digging, sport	26
-----------------------------------	----

Once you get above 26, the risk starts to be high. At a Wet Bulb **Globe** Temperature of 32, the army suspends all training.²⁶⁶ At a Wet Bulb **Globe** Temperature of >39, one cannot survive in the shade for more than a few hours.²⁶⁷

Physical labour becomes difficult to impossible when Wet Bulb Temperature exceeds 31°C.²⁶⁸ At a Wet Bulb Temperature of more than 35, humans cannot survive for more than a few hours, even if they are in the shade, doused in water and have a fan pointed at them.

The lethality of heatwaves depends on the preparedness of the subject population so isn't that reliable a guide to the objective heat stress that people face. For example, the heat wave in Europe in 2003 killed thousands, but the one in 2005 didn't even though temperatures were similar.²⁶⁹

6.2. Heat stress today

The tropics and subtropics and coastal areas are most at risk from heat stress. This is average peak daily Wet Bulb **Globe** Temperature in the hottest month of the year, for 1981-2010:

²⁶⁶ Katharine M. Willett and Steven Sherwood, 'Exceedance of Heat Index Thresholds for 15 Regions under a Warming Climate Using the Wet-Bulb Globe Temperature', *International Journal of Climatology* 32, no. 2 (2012): 161–77, <https://doi.org/10.1002/joc.2257>.

²⁶⁷ "For day-time heat we set the threshold for survivability according to the WBGT that causes core body temperature to rise to 42°C, for an average individual at rest,ii in the shade, for four hours. We estimate this occurs when the daily maximum WBGT is ≥ 40°C." David King et al., 'Climate Change—a Risk Assessment' (Centre for Science Policy, University of Cambridge, 2015), 57, www.csap.cam.ac.uk/projects/climate-change-risk-assessment/.

²⁶⁸ Jonathan R. Buzan and Matthew Huber, 'Moist Heat Stress on a Hotter Earth', *Annual Review of Earth and Planetary Sciences* 48, no. 1 (2020): 623–55, <https://doi.org/10.1146/annurev-earth-053018-060100>.

²⁶⁹ "Echoes of this can be seen in the fact that the European heat wave of 2005 killed very few compared to 2003, despite being just as hot." Jonathan R. Buzan and Matthew Huber, 'Moist Heat Stress on a Hotter Earth', *Annual Review of Earth and Planetary Sciences* 48, no. 1 (2020): 623–55, <https://doi.org/10.1146/annurev-earth-053018-060100>.

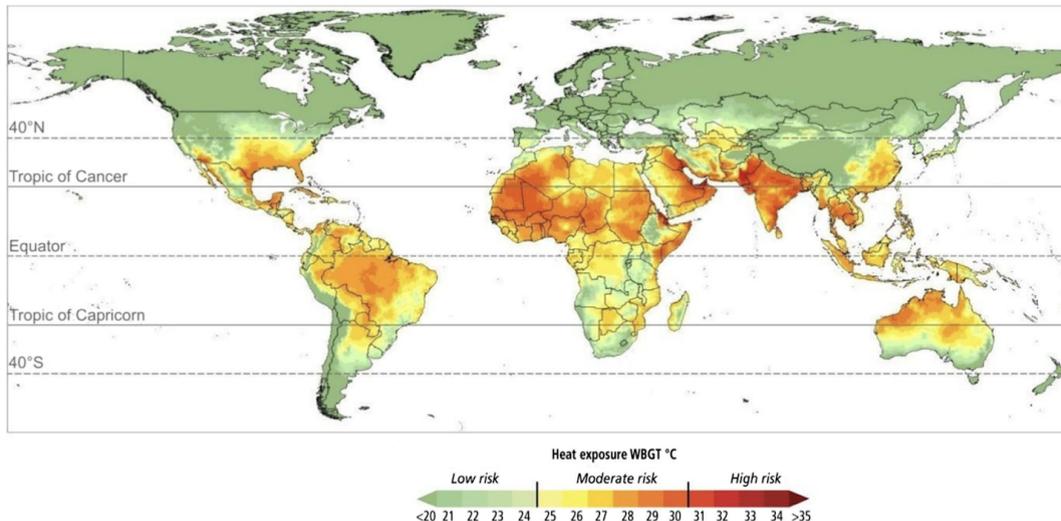
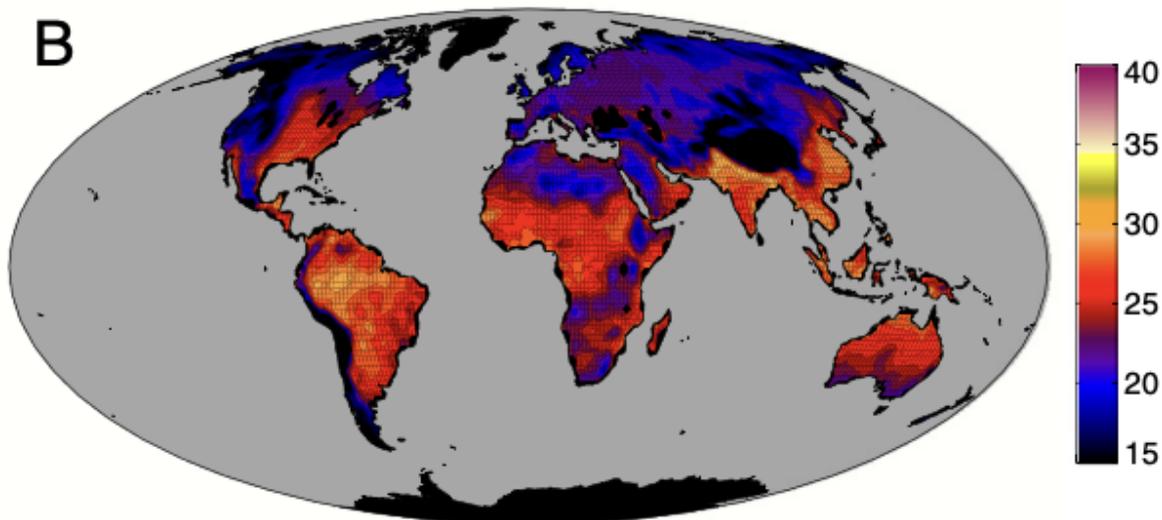


Fig. 2 Average of daily WBGTmax during the locally hottest month in 67,420 grid cells, 30-year mean 1981–2010, CRU data (Similar map was published in the IPCC report, Smith et al. 2014)

Source: Tord Kjellstrom et al., ‘Estimating Population Heat Exposure and Impacts on Working People in Conjunction with Climate Change’, *International Journal of Biometeorology* 62, no. 3 (1 March 2018): 291–306, <https://doi.org/10.1007/s00484-017-1407-0>.

These levels of Wet Bulb **Globe** Temperature would probably obtain for around four hours.²⁷⁰

The chart below shows a map of annual maxima of Wet Bulb Temperature. Recall that a Wet Bulb Temperature of 35 is a hard survival limit.



Source: Steven C. Sherwood and Matthew Huber, ‘An Adaptability Limit to Climate Change Due to Heat Stress’, *Proceedings of the National Academy of Sciences* 107, no. 21 (25 May 2010): fig. 1, <https://doi.org/10.1073/pnas.0913352107>.

²⁷⁰ Tord Kjellstrom, personal correspondence, 28 Jan 2021.

Adjacent nighttime minima of Wet Bulb Temperature are typically within 2–3 °C of the daytime, and adjacent daily maxima are typically within 1°C.²⁷¹ So, within-day and day-to-day changes don't reduce heat stress that much.

Raymond et al (2020) recently showed that weather stations in the UAE and Pakistan have recorded Wet Bulb Temperature above 35 several times,²⁷² and multiple places in South Asia and the Middle East above 30 Wet Bulb Temperature. Numerous other places also suffer very high levels of heat stress.²⁷³

These events do not seem to have been associated with morbidity and mortality,²⁷⁴ which illustrates that thus far people have taken adaptive measures. In the UAE, this would likely include air conditioning, but I am less sure how far people in Pakistan could have taken this step. According to Orlov et al (2020) in South Asia fewer than 10% of households have air conditioning.²⁷⁵ For reference, the average cost of air conditioning for households in the southeastern US is about [\\$525 per year](#).

For the Global Burden of Disease study, Zhao et al (2021) find that between 2000 and 2019, on average there were around 500,000 heat-related deaths and around 4.5 million cold-related deaths each year.²⁷⁶ The literature suggests that the benefits of climate change in reducing cold-related deaths in temperate regions will be outweighed by the increase in heat-related deaths in warmer regions.²⁷⁷

²⁷¹ Sherwood and Huber, 'An Adaptability Limit to Climate Change Due to Heat Stress', 9554.

²⁷² "In presenting boxplots of all TW recordings by month, Fig. S20 makes a similar argument for the two stations with the most TW=35°C readings (Ras Al-Khaimah, UAE, and Jacobabad, Pakistan)." Colin Raymond, Tom Matthews, and Radley M. Horton, 'The Emergence of Heat and Humidity Too Severe for Human Tolerance', *Science Advances* 6, no. 19 (1 May 2020): SI p3, <https://doi.org/10.1126/sciadv.aaw1838>.

²⁷³ "Furthermore, some regions are already experiencing heat stress conditions approaching the upper limits of labour productivity and human survivability (high confidence). These include the Persian Gulf and adjacent land areas, parts of the Indus River Valley, eastern coastal India, Pakistan, north-western India, the shores of the Red Sea, the Gulf of California, the southern Gulf of Mexico, and coastal Venezuela and 27 Guyana (Krakauer et al., 2020);(Li et al., 2020);(Raymond et al., 2020);(Saeed et al., 2021);(Xu et al., 2020)." IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, chap. 7, sec. 7.2.4.1.

²⁷⁴ "Such efforts may also help resolve the reasons for the paucity of reported mortality and morbidity impacts associated with observed near 35°C conditions." Raymond, Matthews, and Horton, 'The Emergence of Heat and Humidity Too Severe for Human Tolerance'.

²⁷⁵ Anton Orlov et al., 'Economic Costs of Heat-Induced Reductions in Worker Productivity Due to Global Warming', *Global Environmental Change* 63 (1 July 2020): Fig. 4, <https://doi.org/10.1016/j.gloenvcha.2020.102087>.

²⁷⁶ Qi Zhao et al., 'Global, Regional, and National Burden of Mortality Associated with Non-Optimal Ambient Temperatures from 2000 to 2019: A Three-Stage Modelling Study', *The Lancet Planetary Health* 5, no. 7 (1 July 2021): Table 1, [https://doi.org/10.1016/S2542-5196\(21\)00081-4](https://doi.org/10.1016/S2542-5196(21)00081-4).

²⁷⁷ "In temperate areas such as northern Europe, east Asia, and Australia, the less intense warming and large decrease in cold-related excess would induce a null or marginally negative net effect, with the net change in 2090–99 compared with 2010–19 ranging from –1.2% (empirical 95% CI –3.6 to 1.4) in Australia to –0.1% (–2.1 to 1.6) in east Asia under the highest emission scenario, although the decreasing trends would reverse during the course of the century. Conversely, warmer regions, such as the central and southern parts of America or Europe, and especially southeast Asia, would experience a sharp surge in heat-related impacts and extremely large net increases, with the net change at the end of the century ranging from 3.0% (–3.0 to 9.3) in Central America to 12.7% (–4.7 to 28.1) in southeast Asia under the highest emission scenario." Antonio Gasparri et al., 'Projections of Temperature-Related Excess Mortality under Climate Change Scenarios', *The Lancet Planetary Health* 1, no. 9 (1 December 2017): e360, [https://doi.org/10.1016/S2542-5196\(17\)30156-0](https://doi.org/10.1016/S2542-5196(17)30156-0).

According to the IPCC, there is evidence from high income countries that the health costs of heat stress are declining over time, due to heat warning systems, increased awareness, and improved quality of life.²⁷⁸ Air conditioning explains up to 20% of this decline.²⁷⁹

6.3. What are high levels of heat stress actually like?

Dhahran in Saudi Arabia is one of the most heat stressed inhabited places in the world and regularly has Wet Bulb Temperature above 30 each year since 1970.²⁸⁰ This is a description sent to the [Washington Post](#) from someone who lived there:

“When the winds come off the Persian Gulf you just can’t imagine how awful it gets. On the hottest and most humid days, you’d walk outside and it felt immediately like someone pressed a hot wet towel, like you sometimes get on airplanes, over your entire head. I wear glasses, and they’d immediately fog up. You sweat instantly. People just avoid being outside in any way they can. In the summers, my friends and I would become nocturnal as a way to beat the heat. Crime is basically non-existent, so my parents didn’t worry about us being out all night. I’d usually have breakfast with my dad and then sleep through the heat of the day, waking up when he got home from work. At night it was still stifling, but the edge was off.

Air conditioning is everywhere. You can trace the population explosion in the country directly to the advent of air conditioning – it allowed people to settle down and stop living the nomadic life that was common into the middle of the 20th century. We lived on a compound for employees of the Saudi national oil company, and they treated air conditioning repair like ambulances or fire trucks – they had crews on 24-hour call, and you could have them dispatched at a moment’s notice by calling the special air conditioning emergency hotline. In the summer, the air-conditioned school buses would stop outside every individual kid’s house, so they didn’t have to wait at a stop and could stay in the AC. Off the compound, air conditioning is still common, even for the poorest migrant workers there. Shopping was done in huge air-conditioned malls. The great open-air souks operate in the winter or very early in the morning on summer weekends”

Friends in Dubai have told me that in summer, people move between different air conditioned environments in their car/office/home.

²⁷⁸ “Several lines of evidence point to a possible decrease in population sensitivity to heat, albeit mainly for high57 income countries (high confidence), arising from the implementation of heat warning systems, increased awareness, and improved quality of life” IPCC, *Climate Change 2022: Impacts*, Ch. 7, sec. 7.2.4.1.

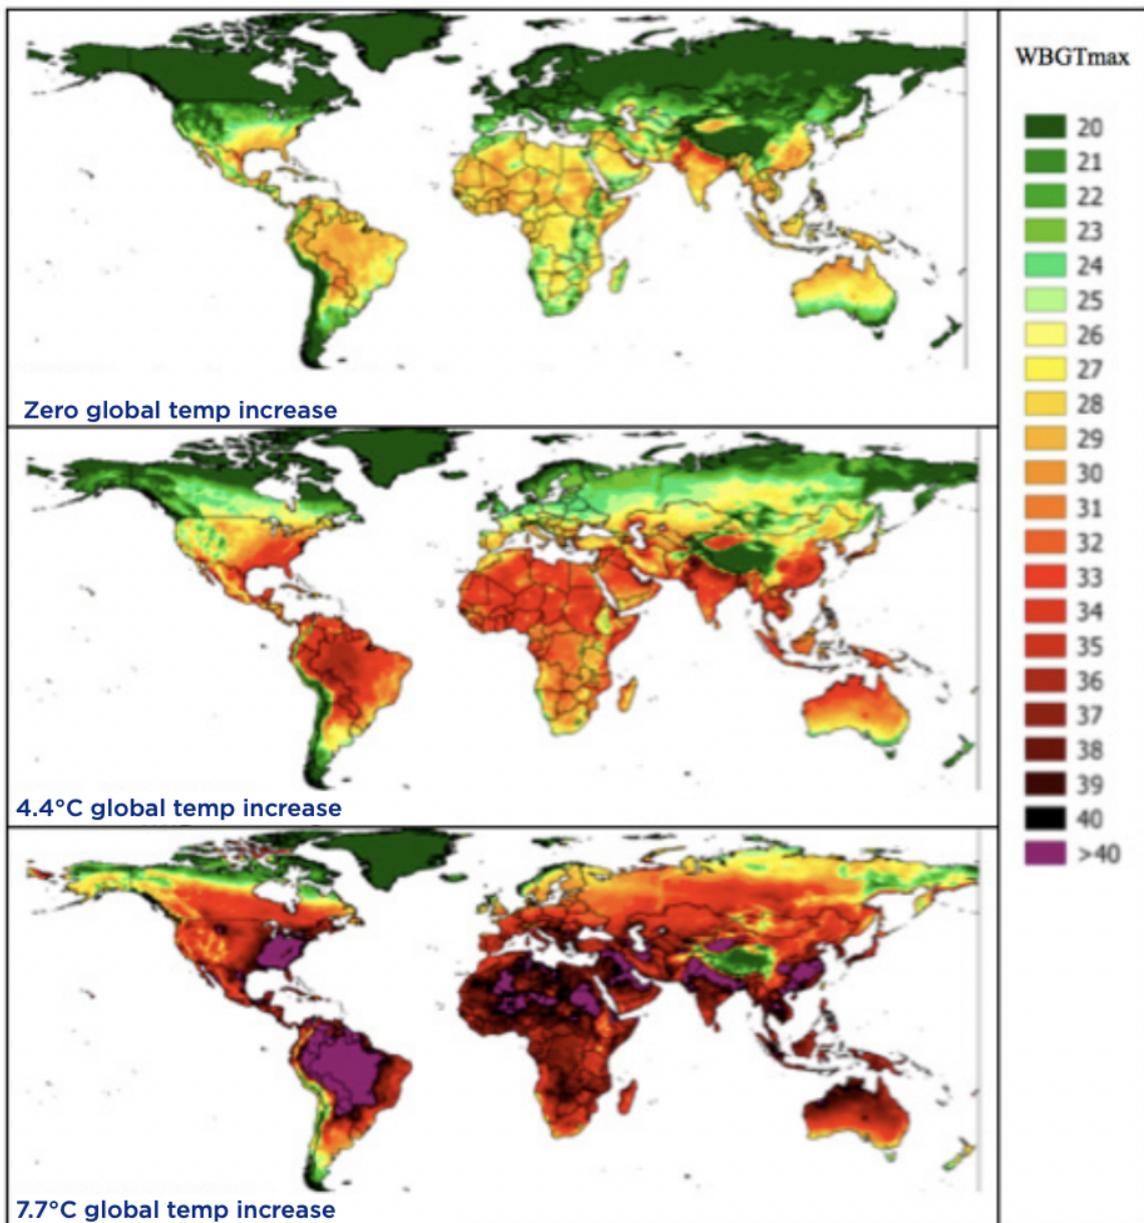
²⁷⁹ “Although there is a paucity of global 46 level studies of the effectiveness of air conditioning for reducing heat-related mortality, a recent assessment 47 indicates increases in air conditioning explains only part of the observed reduction in heat-related excess 48 deaths, amounting to 16.7% in Canada, 20.0% in Japan, 14.3% in Spain and 16.7% in the US (Sera et al., 49 2020).” IPCC, *Climate Change 2022: Impacts*, Ch. 7, sec. 7.2.4.1.

²⁸⁰ Jeremy S. Pal and Elfatih A. B. Eltahir, ‘Future Temperature in Southwest Asia Projected to Exceed a Threshold for Human Adaptability’, *Nature Climate Change* 6, no. 2 (February 2016): fig. 2, <https://doi.org/10.1038/nclimate2833>.

6.4. Impacts

Rising heat stress would increase heat-related deaths and reduce outdoor labour capacity for regions that cannot afford to adapt. It is important to stress that most future population growth will be in the tropics, which will be hardest hit by rising heat stress.

The map below shows average daily peak Wet Bulb **Globe** Temperature during the hottest month of the year for today, for 4.4°C and 7.7°C of warming.



Source: King et al., 'Climate Change—a Risk Assessment', 62.

For reference, orange is above the safe limit for walking for acclimatised people, according to international standards. In the summer months, outdoor activity would be very difficult in

the tropics. At 6°C of warming, New York City would become more heat stressed than Bahrain is today.²⁸¹

This could have very bad effects on labour capacity for countries with limited adaptive capacity. Using standard safety guidelines, Buzan and Huber (2020) calculate the reductions in population-weighted labour capacity at different levels of warming assuming no air conditioning or other adaptation, and that people do not migrate:

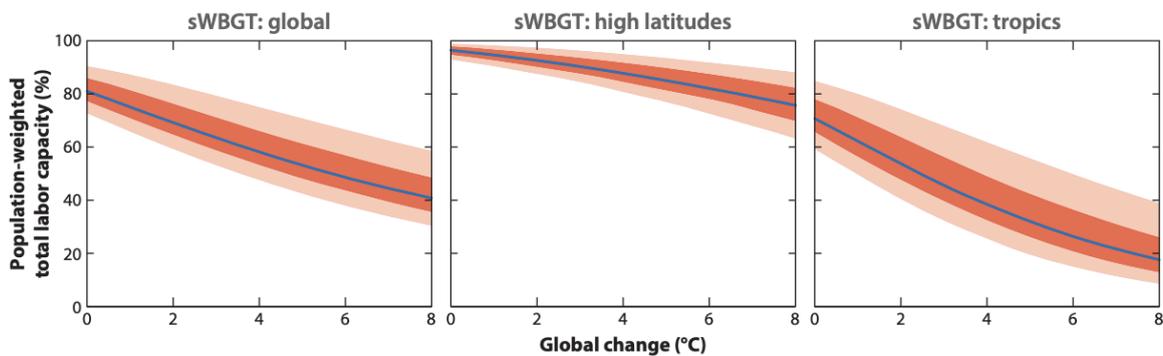


Figure 10

Population weighted total labor capacity. The CMIP5 ensemble is represented by the median (*blue line*), 50% (*red swath*), and 80% (*pink swath*) confidence intervals. The relative impacts on labor are shown at global (57°S to 57°N), high latitude (outside of 30°S to 30°N), and tropic (30°S to 30°N) regions. Abbreviations: CMIP, Coupled Model Intercomparison Project; sWBGT, simplified wet bulb globe temperature.

Source: Buzan and Huber, 'Moist Heat Stress on a Hotter Earth'.

As this shows, while the high latitudes emerge relatively unscathed, the tropics are especially badly affected by rising heat stress, with labour capacity falling to 40% of its capacity at 5°C and 20% of its capacity at 8°C.

This estimate is a pessimistic upper bound on the effect because it assumes no adaptation and no migration. I discuss the prospects for adaptation below.

Turning to hard survivability limits, the chart below shows how heat stress will change in the Middle East and Persian Gulf, one of the world's most heat stressed regions. This shows annual maximum Wet Bulb Temperature on 2.7°C of warming (green line) and 4.4°C (red line).

²⁸¹ John P. Dunne, Ronald J. Stouffer, and Jasmin G. John, 'Reductions in Labour Capacity from Heat Stress under Climate Warming', *Nature Climate Change* 3, no. 6 (June 2013): 563–66, <https://doi.org/10.1038/nclimate1827>, fig. 1.

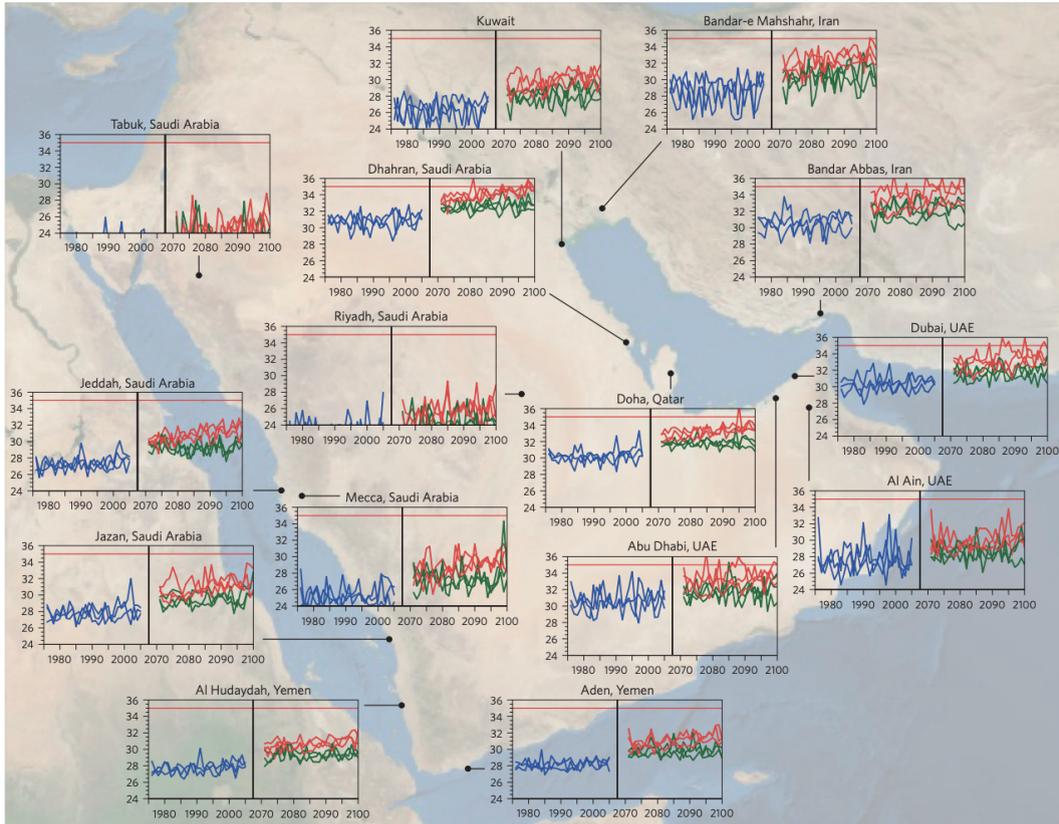
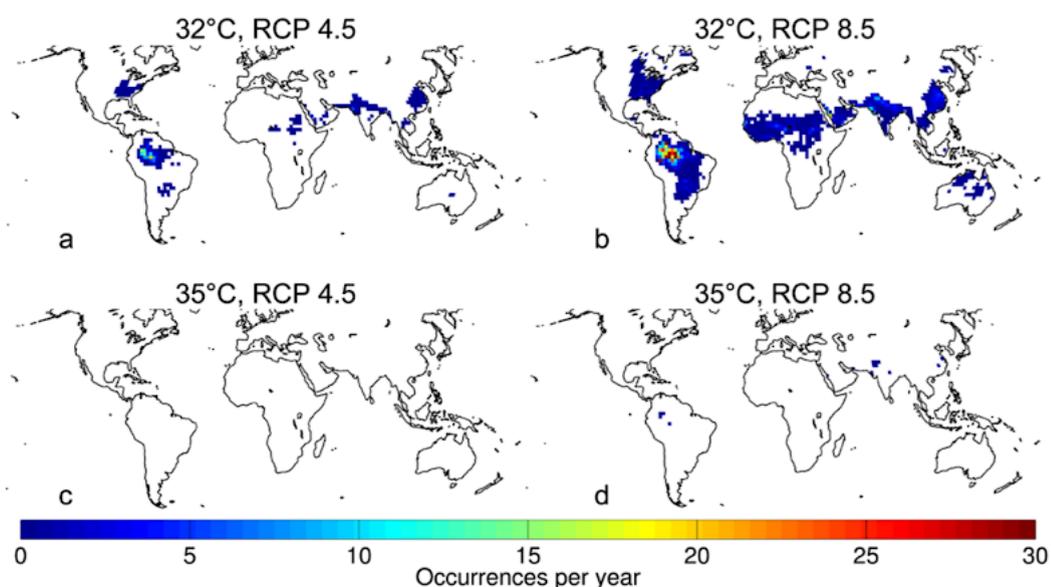


Figure 2 | Time series of the annual maximum TW_{max} for each ensemble member and GHG scenario. Blue, green and red lines represent the historical (1976–2005), RCP4.5 (2071–2100) and RCP8.5 (2071–2100) scenarios, respectively. TW_{max} is the maximum daily value averaged over a 6-h window. The background image was obtained from NASA Visible Earth.

Source: Jeremy S. Pal and Elfatih A. B. Eltahir, 'Future Temperature in Southwest Asia Projected to Exceed a Threshold for Human Adaptability', *Nature Climate Change* 6, no. 2 (February 2016): fig. 2, <https://doi.org/10.1038/nclimate2833>.

The map below shows the number of days that people are exposed to Wet Bulb Temperature above different levels for 2.7°C of warming (RCP4.5) and 4.4°C of warming (RCP8.5).



Supplementary Figure 5: Multi-GCM mean number of days in 2070 – 2080 with wet bulb temperatures above 32°C (top row) and 35°C (bottom row). Left panels show results under RCP 4.5 and right panels under RCP 8.5. Wet bulb temperatures above 35°C are limited to small geographic areas, even under RCP 8.5, but some of these regions – in particular northeastern India and eastern China – are densely populated. RCP 4.5 completely avoids wet bulb temperatures of 35°C through 2080.

Source: Coffel, Horton, and Sherbinin, 'Temperature and Humidity Based Projections of a Rapid Rise in Global Heat Stress Exposure during the 21st Century', fig. S5.

Recall that outdoor labour is difficult or impossible once Wet Bulb Temperature passes 31°C. As this shows, for 2.7°C of warming, regions in the tropics and subtropics experience Wet Bulb Temperature above 32°C for around 1-10 days each year. This is close to the most heat stressed places in the Persian Gulf today. For 4.4°C of warming, a much larger number of people would be exposed to 1-10 days of Wet Bulb Temperature above 32 each year.

These temperature extremes will persist across much of the summer in the places affected, as temperatures don't drop much either side of the hottest month. For example, May is the hottest month in Bihar in India at 38°C (dry bulb i.e. measured by a normal thermometer) and April and June are around 37°C.²⁸² So, we would be looking at regularly exceeding a Wet Bulb Temperature of 30 in the summer months in these places.

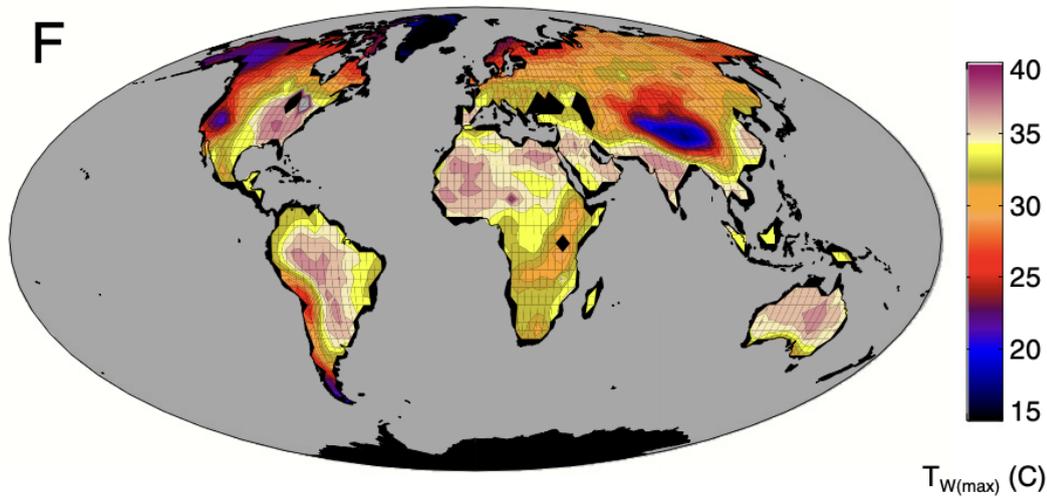
A good summary of the recent literature on heat stress associated with 4°C of warming is in Coffel et al (2017),²⁸³ and there are various recent papers by Eltahir and others on heat stress in specific regions.²⁸⁴

²⁸² Data here is from [Accuweather](#)

²⁸³ Coffel, Horton, and Sherbinin, 'Temperature and Humidity Based Projections of a Rapid Rise in Global Heat Stress Exposure during the 21st Century'.

²⁸⁴ Pal and Eltahir, 'Future Temperature in Southwest Asia Projected to Exceed a Threshold for Human Adaptability'; Eun-Soon Im, Jeremy S. Pal, and Elfatih A. B. Eltahir, 'Deadly Heat Waves Projected in the Densely Populated Agricultural Regions of South Asia', *Science Advances* 3, no. 8 (1 August 2017): e1603322, <https://doi.org/10.1126/sciadv.1603322>; Suchul Kang and Elfatih A. B.

The map below from the classic Sherwood and Huber (2010) shows annual maximum Wet Bulb Temperature for 11°C of warming above pre-industrial:



Source: Sherwood and Huber, 'An Adaptability Limit to Climate Change Due to Heat Stress', fig. 1.

In this extreme worst-case scenario. Heat stress levels would pass lethal limits at some point in the year for the majority of the global population, as people are currently distributed. People will still be able to survive at higher latitudes and altitudes, and people in the tropics could survive in air conditioned environments. Humanity would be much diminished, though we would survive.

6.4.1. The scope for adaptation

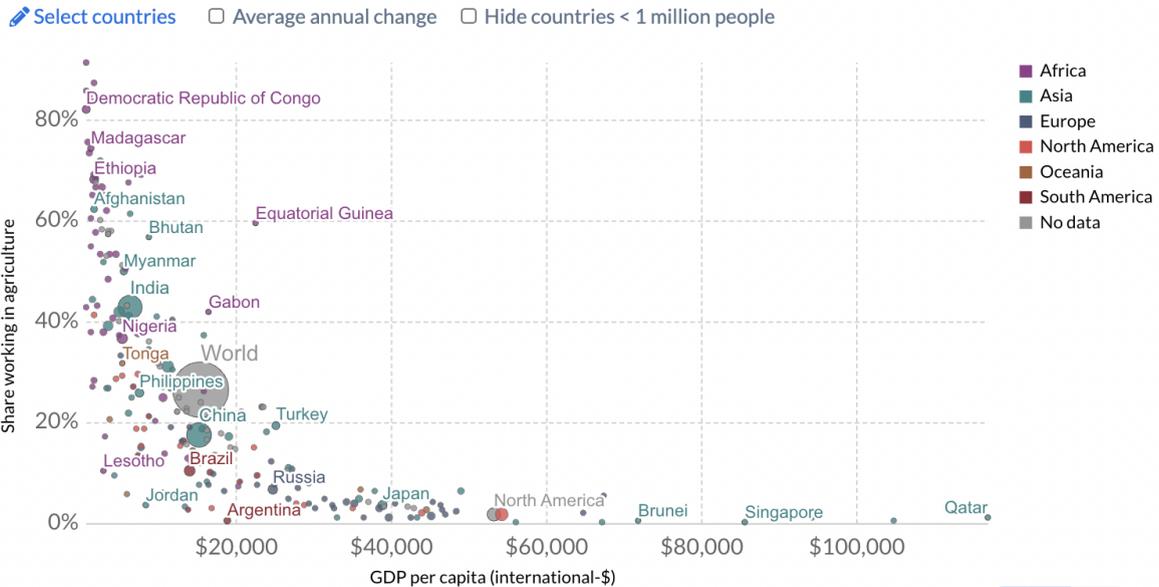
Despite rising levels of heat stress, in rich countries at least, heat-related deaths are declining. As countries get richer, we should expect the effects of heat stress to become less bad, other things being equal.

One important factor is that in growing regions, a smaller fraction of the global population will work outdoors in agriculture. Growing economies will have income per head in excess of \$20,000 per person by 2100, which suggests that agricultural employees will constitute less than 10% of the workforce, compared to a global average today of around 25%.

Eltahir, 'North China Plain Threatened by Deadly Heatwaves Due to Climate Change and Irrigation', *Nature Communications* 9, no. 1 (31 July 2018): 2894, <https://doi.org/10.1038/s41467-018-05252-y>.

Employment in agriculture vs GDP per capita, 2017

Share of persons of working age who were engaged in any activity to produce goods or provide services for pay or profit in the agriculture sector (agriculture, hunting, forestry and fishing).



Orlov et al (2020) argue that there is scope for outdoor workers in construction and agriculture to adapt by using mechanisation: it is easier to farm with a tractor and machinery than by hand. This could slightly reduce the energy requirements of outdoor labour.

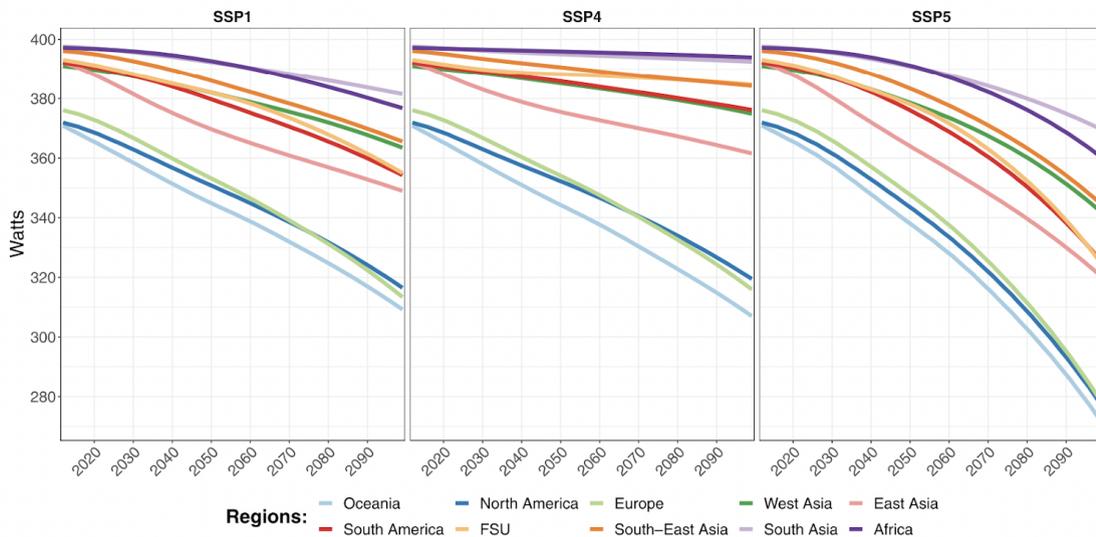


Fig. 5. Calibrated relationship between economic growth and work intensity in agriculture and construction.

Source: Orlov et al., 'Economic Costs of Heat-Induced Reductions in Worker Productivity Due to Global Warming'.

Another reason that richer regions can better adapt to heat stress is that they can afford air conditioning and other adaptive measures. Orlov et al (2020) project air conditioning trends in different regions on different Shared Socioeconomic Pathways.

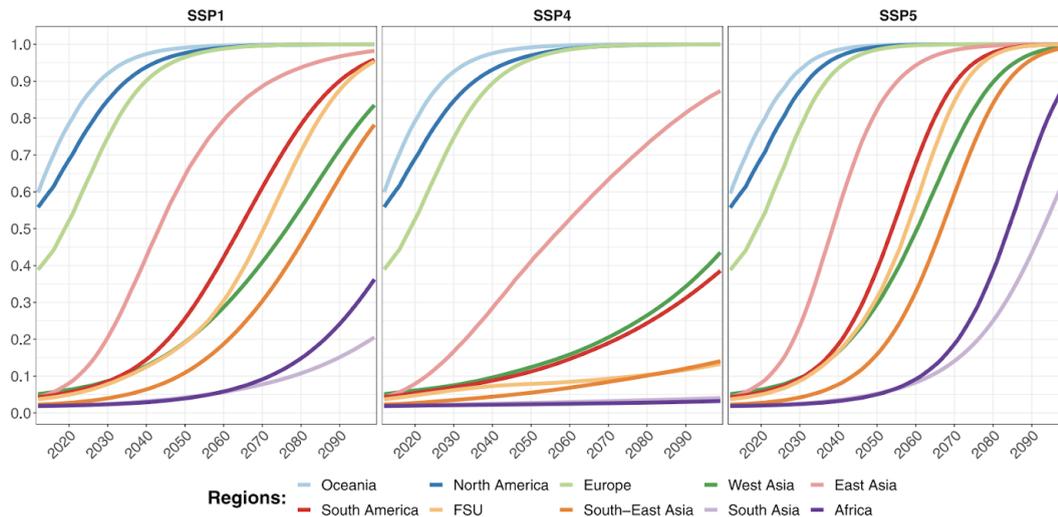


Fig. 4. Penetration rates of air conditioners by region and SSP.

Source: Orlov et al., 'Economic Costs of Heat-Induced Reductions in Worker Productivity Due to Global Warming'.

There is a large range in the penetration of air conditioning in the poorest regions - South Asia and Sub-Saharan Africa - depending on the SSP. On the high growth SSP5, penetration is above 60%, whereas on the unequal SSP4, it is close to zero.

Although there is a fair bit of scope for adaptation in rich and/or growing regions, to say the least, the $>4^{\circ}\text{C}$ world is not a compelling vision of the future. Billions of people in the tropics and subtropics would be unable to do much outside during the summer months, including things like jogging or even walking for a few hours. People would have to move between different air conditioned pods, as people do in Dubai in summer. For countries that are still reliant on agriculture and outdoor labour in 2100, this would be very damaging.

6.4.2. Death estimates

Bressler et al (2021) quantifies the effect that climate change will have on global all-cause mortality with and without adaptation

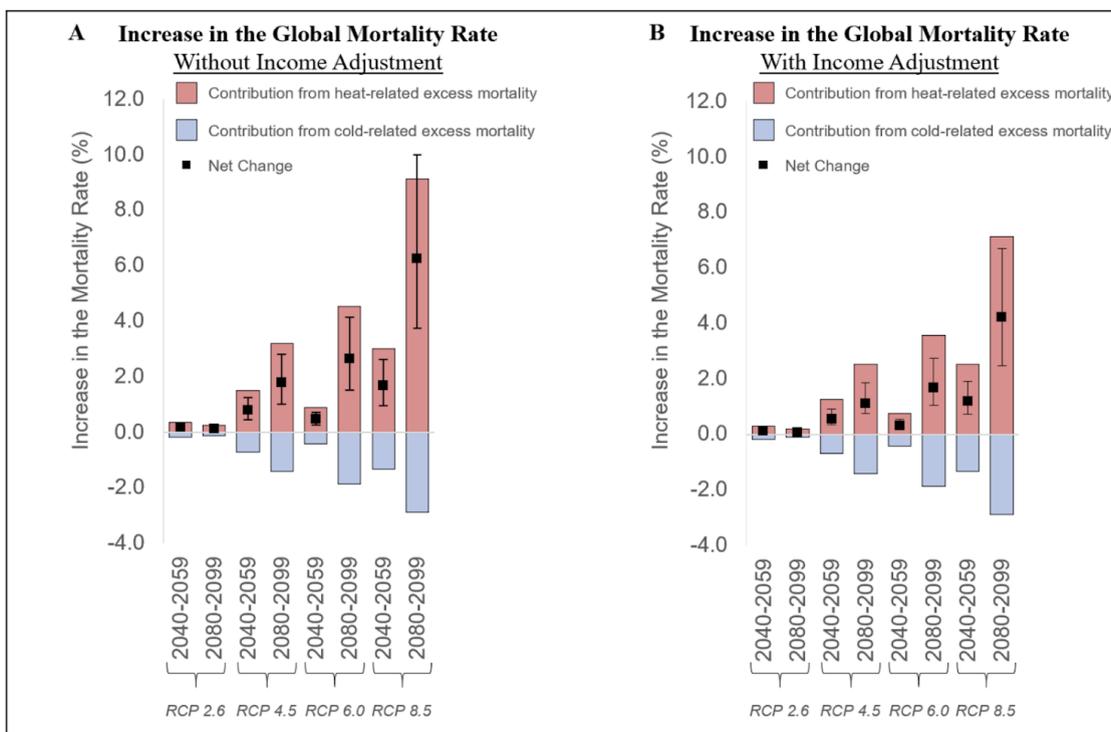


Figure 5. Increase in the global mortality rate by scenario. The graphs show the contribution of global warming towards the mortality rate through its effect on temperature related mortality, without (Panel A) and with (Panel B) including the estimated protective effects of higher incomes, projected using SSP3. The black error bars show 95% CIs for the net mortality impact.

Source: R. Daniel Bressler et al., 'Estimates of Country Level Temperature-Related Mortality Damage Functions', *Scientific Reports* 11, no. 1 (13 October 2021): 20282, <https://doi.org/10.1038/s41598-021-99156-5>.

With income-related adaptation, on RCP4.5, which implies 2.7°C of warming, the global mortality rate would increase by around 1%. On the UNs median projections, deaths would increase by around 1 million per year by 2100. On RCP8.5, which implies around 4.4°C of warming, the global mortality rate would increase by 4.2%, which translates to 5 million extra deaths per year, on net. (Calculations are [here](#)).

There are two reasons to think that this estimate is overly pessimistic. Firstly, it excludes some adaptive measures.

“It is possible that other forms of adaptation and technological changes, other than through income-based adjustments, might modify the temperature-mortality relationship over time. For example, provision of public heat alert systems, improved preparation of the medical system for heat-related diseases, or people learning to avoid activity during the hottest parts of the day might all reduce the adverse effects of extreme heat over time. Several studies reviewed by Arbuthnott et al. show evidence of decreasing sensitivity of heat-related mortality over time. These effects are not included in the estimates given here.”

The main income-based adaptation that Bressler et al (2021) mention is air conditioning.²⁸⁵ According to the IPCC, air conditioning only explains 15-20% of the decline in heat-related deaths seen in high-income countries. This suggests that by focusing only on income-based adaptation, Bressler et al (2021) may underestimate the total effects of adaptation by up to a factor of five. For 4.4°C of warming, the adaptation measures considered by Bressler et al reduce the increase in the death rate from 6.2% to 4.2%. If this 2 percentage point reduction in the death rate is indeed underestimated by a factor of 5, then the heat-related effects of warming would be eliminated.

Secondly, Bressler et al (2021) considers the impacts of heat stress conditional on pessimistic assumptions about socioeconomic development. They consider the impacts of RCP4.5 and RCP8.5 on Shared Socioeconomic Pathway 3. SSP3 is the lowest growth SSP, with particularly sluggish growth for Africa and parts of Asia. The predicted growth for Asia seems overly pessimistic, which is important for future heat stress projections.

As discussed in Chapter 1, RCP8.5 is only possible on SSP5. The ‘current policy’ baseline on SSP3 is RCP7. Moreover, as I said in Chapter 1, this baseline scenario seems far too pessimistic about likely emissions on SSP3 and on current policy. This being said, Bressler et al (2021) only measures the *most likely* level of warming conditional on a given emissions pathway: for RCP8.5, this is around 4.4°C. But warming might be higher than we expect. On RCP7, the chance of 4.5°C is around 1 in 20.²⁸⁶

Overall, the 5 million deaths estimate should be taken as a probably pessimistic estimate of the most likely level of deaths due to heat stress at 4.4°C of warming.

6.5. The future human climate niche

‘The Future of the Human Climate Niche’ is an interesting 2020 paper by Xu et al which explores the climatic environments that people live in today and how they will change with future warming.²⁸⁷ Xu et al only explore the potential impact of temperature, and exclude other effects of climate change, such as changes in water stress. Xu et al discuss the effects of temperature change on heat stress, economic productivity and general livability.

Xu et al note that “human populations have resided in the same narrow part of the climatic envelope available on the globe, characterized by a major mode around ~11 °C to 15 °C

²⁸⁵ See for example “Finally, Model 4 adds an additional interaction term with per-capita income, reflecting the ability of individuals and groups to make investments that mitigate the negative mortality effect of heat, such as installing air conditioning”; “We also see a negative interaction with log (GDPPCc) (i.e. a negative β_4 coefficient) indicating that richer countries can ameliorate some of the damages associated with higher temperatures. This could well be associated with air conditioning penetration, which several studies have shown to be strongly associated with higher incomes, particularly in warmer, middle-income countries”; “When the reduced sensitivity to heat associated with rising incomes, such as greater ability to invest in air conditioning, is accounted for, the expected end-of-century increase in the global mortality rate is 1.1% [95% CI 0.4–1.9%] in RCP 4.5 and 4.2% [95% CI 1.8–6.7%] in RCP 8.5.” R. Daniel Bressler et al., ‘Estimates of Country Level Temperature-Related Mortality Damage Functions’, *Scientific Reports* 11, no. 1 (13 October 2021): 20282, <https://doi.org/10.1038/s41598-021-99156-5>.

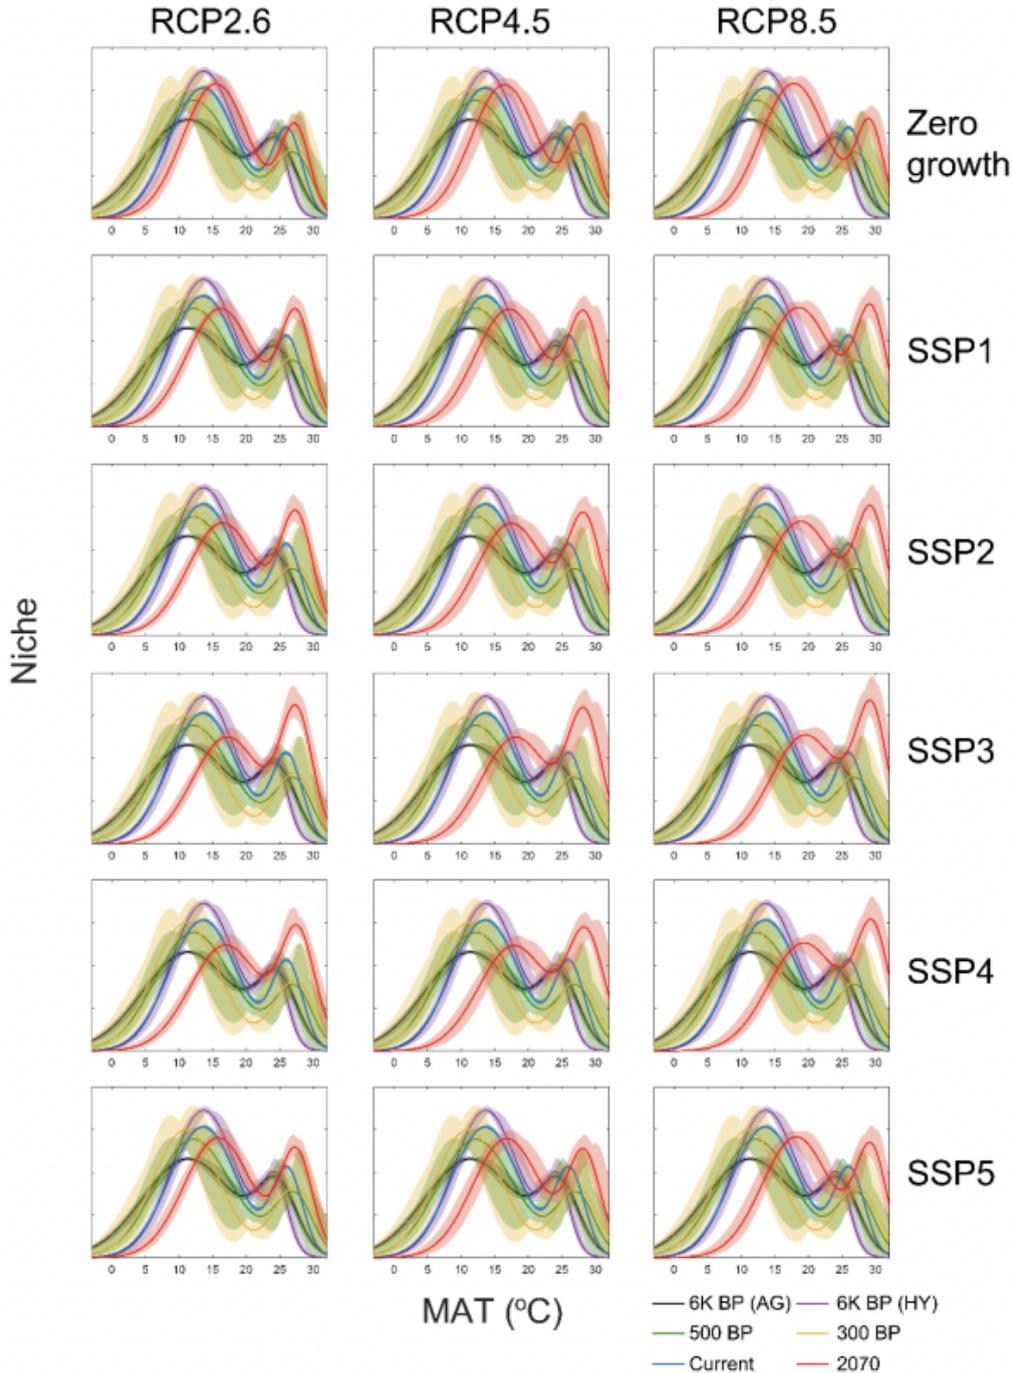
²⁸⁶ See Chapter 1.

²⁸⁷ Chi Xu et al., ‘Future of the Human Climate Niche’, *Proceedings of the National Academy of Sciences* 117, no. 21 (26 May 2020): 11350–55, <https://doi.org/10.1073/pnas.1910114117>.

mean annual temperature". On the most pessimistic scenario to 2070, which assumes SSP3 and RCP8.5, and that people do not migrate, by 2070, due to climate change and population growth in hotter regions, one third of the global population will experience an average annual temperature of more than 29°C, a situation found in the present climate only in 0.8% of the global land surface, mostly concentrated in the Sahara, but in 2070 projected to cover 19% of the global land.

The chart below shows the future human climate niche on different socioeconomic and emissions scenarios:

Figure S7. Past and current human niche (fitted by double Gaussian models) in terms of mean annual temperature (MAT), contrasted to the projected situation in 2070 (red). Bands represent 5th and 95th percentiles of the ensemble of climate and population reconstructions. Different scenarios of climate (RCP2.6, RCP4.5 and RCP8.5) and population growth (zero growth, and SSP1-5) were considered. Bands represent 5th and 95th percentiles of the ensemble of climate and population reconstructions. Reconstructed population data based on the HYDE 3.1 (HY) and ArchaeoGlobe (AG) database were used for 6Ky BP.



Source: Xu et al., 'Future of the Human Climate Niche', SI Fig. S7.

For reference, on RCP2.6, warming at 2070 would be around 1.7°C. On RCP4.5, it would be around 2.3°C, and on RCP8.5, it would be around 3.4°C.

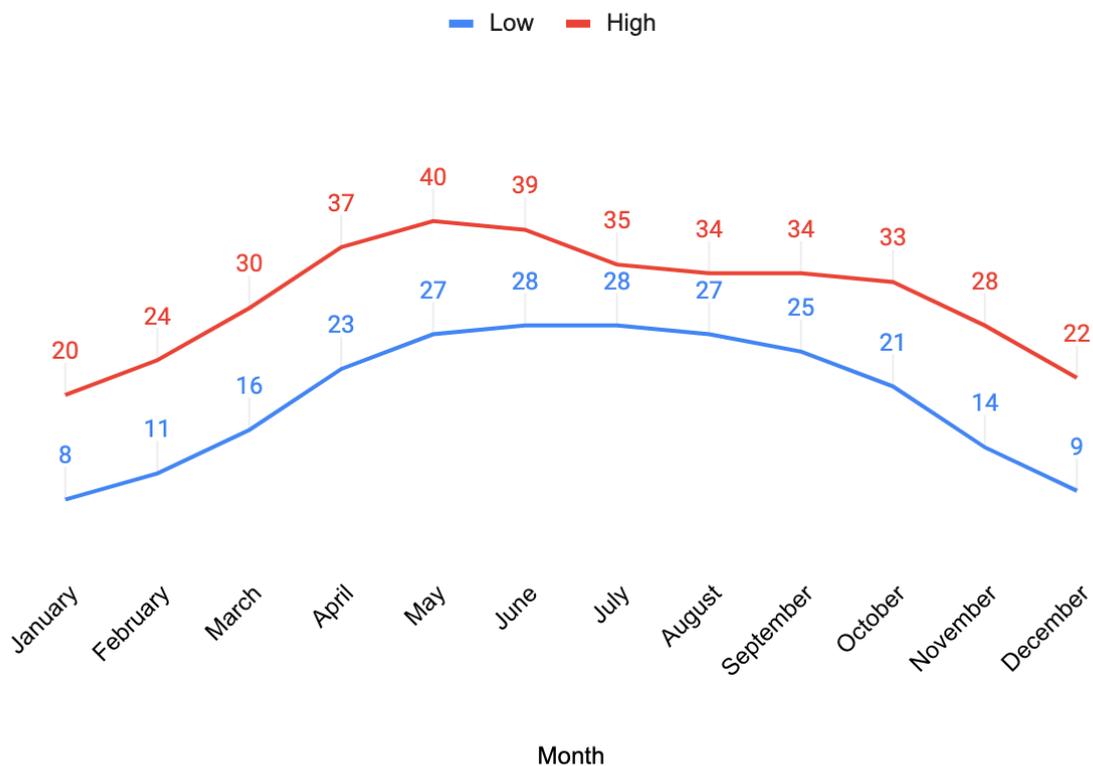
As Xu et al note, their model also assumes no migration, but this is a likely adaptive response to increased heat stress.

One can examine the effect that population growth in tropical regions has on these projections by comparing the 'zero [population] growth' scenario to the SSPs. The SSP with the highest population growth is SSP3. On RCP8.5, the increase in the number of people experiencing a mean annual temperature above 25°C is several times smaller in the world with no population growth, compared to the world with SSP3-levels of population growth. Global warming and population growth in hot places work in tandem to increase the population exposed to higher temperatures.

It is instructive to explore how mean annual temperature experienced by humans will change in the future, but we can get an even richer picture by exploring more of the temperature range experienced by humans over the course of a year. People do not live the whole year at an average temperature. Rather they live through large diurnal and seasonal temperature changes. One way to guide intuitions about the effects of climate change is to think of it as *shifting up the temperature distribution in different locations* across the world. It must again be stressed that the shift will not be uniform across the world.

For example, the annual average temperature for people in India is around 25°C. But residents of Delhi live through large temperature changes over the course of a year:

Temperature in New Delhi

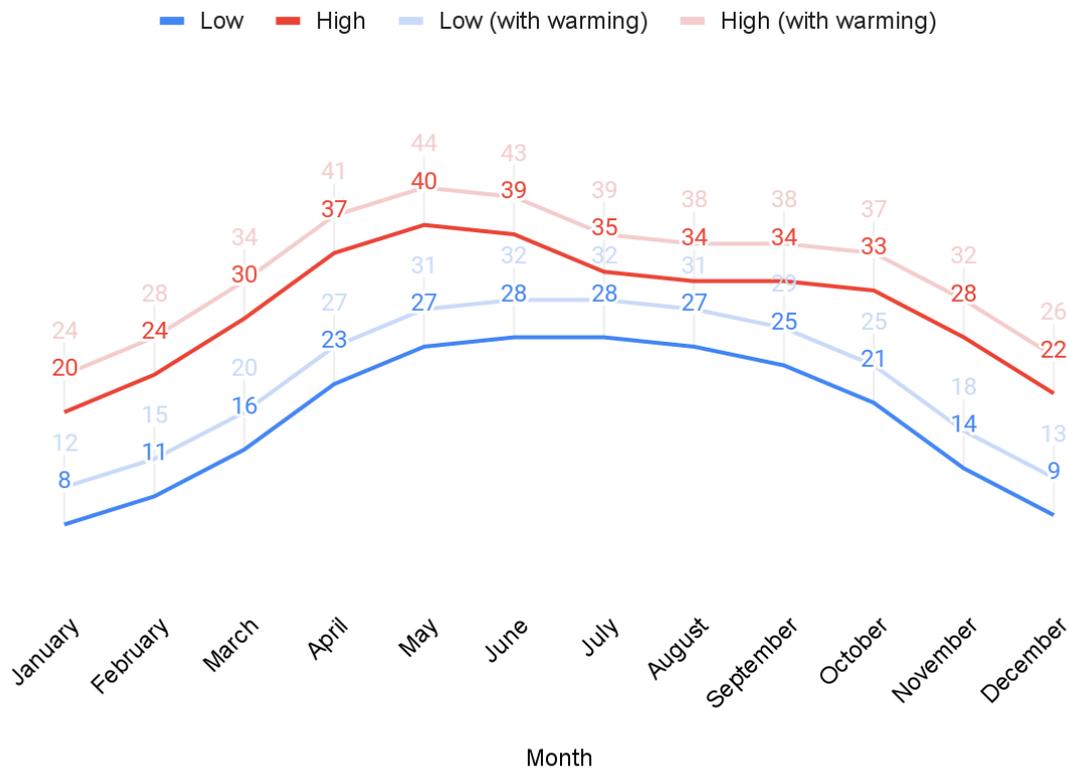


Source: Google weather

In January, Delhi residents live through diurnal temperature variation of more than 10°C - from monthly lows of 8°C to highs of 21°C. In the transition from winter to summer, people will live through a temperature change of 22°C: from the winter low of 8°C to the summer high of 40°C.

Climate change shifts these distributions up on the y-axis. For instance, here are high and low temperatures in Delhi assuming uniform 4°C of local warming across the year:

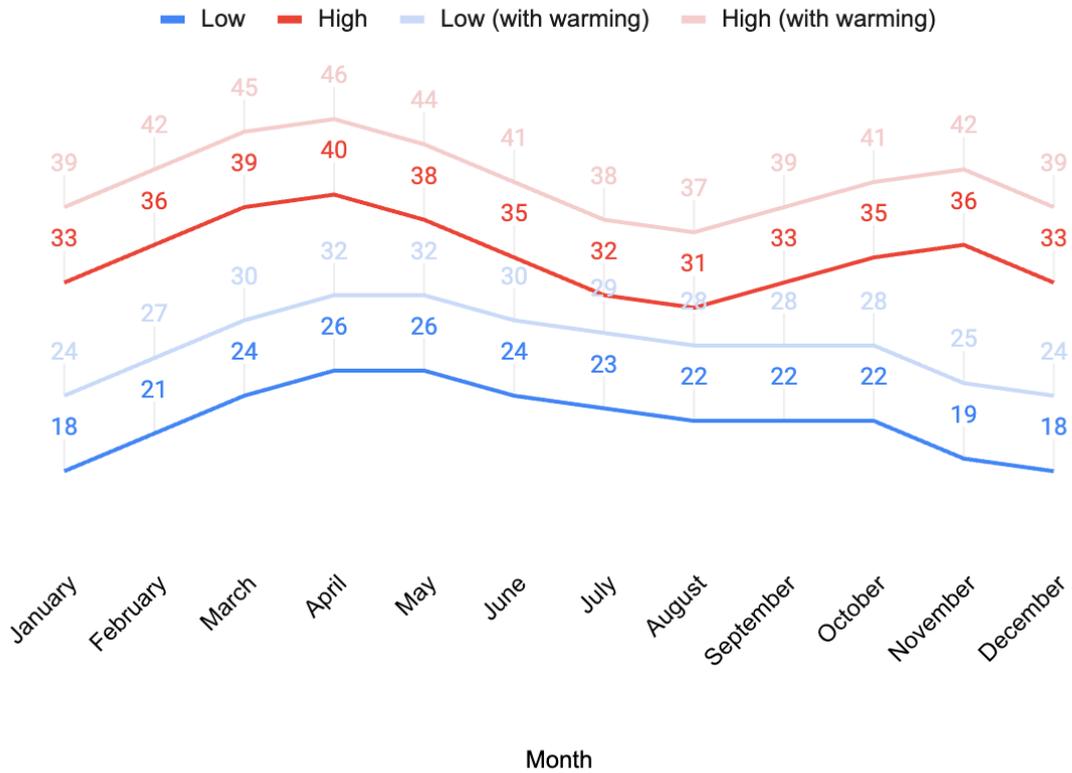
Temperature in New Delhi



As this shows, warming exposes Delhi residents to a new higher peak summer temperature. However, for the vast majority of the year, people live within the same climate envelope that they did before. Except in May and June, the temperatures experienced by Delhi residents are all within the pre-warming temperature envelope. Even in May and June, for the majority of the time, people will be within the pre-warming envelope.

Global warming is not uniform. For Delhi residents, global average surface warming will in fact be fairly close to local Delhi warming. For others, such as people in the Middle East, in central Africa and northern Europe, the local temperature will be higher, according to climate models. For instance, in Bamako, the capital of Mali, local warming will be far higher than global average surface warming; 4°C of global warming would translate into closer to 6°C of local warming. Thus, the temperature envelope would change as follows:

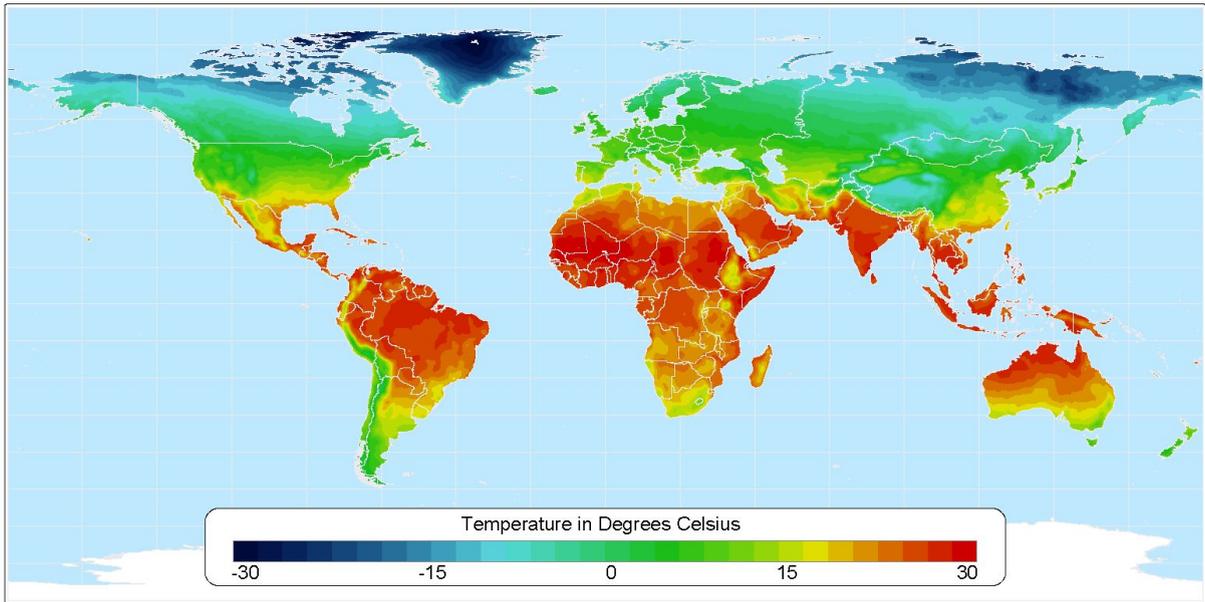
Temperature in Bamako



In Bamako, people would still spend the majority of their year within the current temperature envelope, but peak mean daily temperatures would be higher than current peak mean daily temperatures for around 7 months of the year.

As I have said, this does not provide the complete picture because it excludes many avenues of climate impact, such as droughts, fires, rising sea level and so on. But I do think this is a better way to guide intuition about climate impacts than focusing on mean annual temperatures.

Annual average temperatures vary substantially across heavily populated regions:



Data taken from: CRU 0.5 Degree Dataset (New, et al.)

Atlas of the Biosphere

Center for Sustainability and the Global Environment
University of Wisconsin - Madison

With 4°C of warming, London would have a temperature similar to Madrid and Madrid would have a temperature similar to Shenzhen. Although this would be very bad, it is hard to see how this could completely destroy modern industrial civilisation: advanced economies already thrive at these annual average temperatures.

7. Sea level rise

Due to global warming, sea levels have risen by about 20cm on average since 1900.²⁸⁸ The effect of CO₂ emissions on sea level over the next few millennia brings home the importance of taking a longterm perspective: sea level rise would be Earth-changing over the coming millennia, even on relatively modest emissions scenarios.

Sea level rise threatens the world's coasts through a range of impacts including²⁸⁹ i) permanent submergence of land by mean sea levels or mean high tides; ii) more frequent or intense coastal flooding; iii) enhanced coastal erosion; loss, degradation, and change of coastal ecosystems; iv) salinization of soils and of ground and surface water; and v) impeded drainage of natural (e.g. rivers) and artificial (e.g. sewage) water systems. These biophysical impacts will in turn have socioeconomic impacts on coastal residents and their livelihoods, such as flood damage to buildings, disruption and relocation of economic activities, migration, and degraded coastal agriculture. Here, I focus mainly on permanent submergence and temporary flooding as these two impacts are the ones best studied.

7.1. Past trends in relative sea level rise and flooding

It is important to note that relative sea level rise is what matters for impacts, and this depends on both global warming-induced sea level rise as well as on local vertical land movement in terms of either land uplift or subsidence.

7.1.2. Vertical land movement

Subsidence is a major problem for many coastal cities around the world, and is mainly man-made.²⁹⁰ Due to subsidence, many cities have experienced very large (>1m) relative sea level rise over the 20th century. The problem is especially bad in Asia.²⁹¹

This chart shows subsidence since 1900 in various cities.

²⁸⁸ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), p. 5.

²⁸⁹ IPCC, 'Special Report on the Ocean and Cryosphere in a Changing Climate'.

²⁹⁰ "Natural subsidence, mainly due to the compaction of young sediments in deltas, is widespread and noteworthy. However, the most rapid rates of subsidence are human-induced. These are caused by accelerated compaction primarily due to withdrawal of underground fluids including groundwater, oil and gas, as well as drainage of organic soils... these processes are marked in many of the world's deltas and are often compounded by both local flood defences within the delta and upstream dams, which collectively reduce the sediment supply that maintains these sedimentary landforms. Sand extraction and mining can exacerbate this loss of sediment supply" Robert J. Nicholls et al., 'A Global Analysis of Subsidence, Relative Sea-Level Change and Coastal Flood Exposure', *Nature Climate Change*, 8 March 2021, 1–5, <https://doi.org/10.1038/s41558-021-00993-z>.

²⁹¹ "Cumulatively, human effects on subsidence are at their largest in some coastal cities located on deltas and alluvial plains: a net subsidence of more than 4m has occurred during the twentieth century in parts of Tokyo, and 2 to 3m in Shanghai, Bangkok, Jakarta and New Orleans. Many deltas and subsiding cities are in Asia, and the World Bank has recognized that subsidence could be as influential as climate-induced SLR in parts of coastal Asia over the twenty-first century" Nicholls et al.

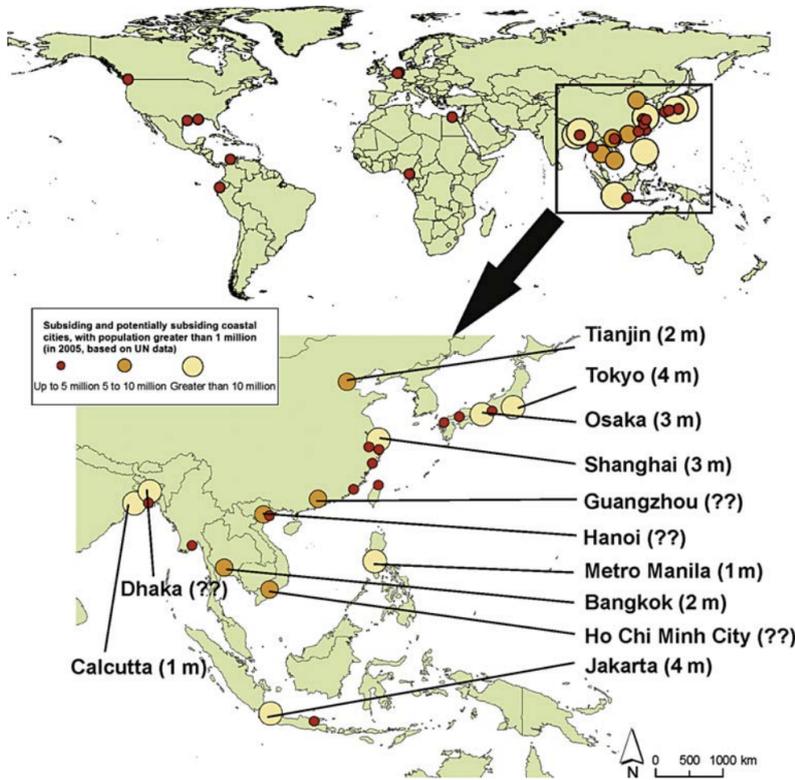


FIGURE 2.3 Subsiding and potentially subsiding coastal cities due to human influence. The maximum observed subsidence (in meters) is shown for cities with populations exceeding 5 million people, where known. The maximum subsidence is reported as data on average subsidence is not available. *Reproduced from Nicholls, R.J., 2014. Adapting to sea level rise. In J.T. Ellis, D.J. Sherman (Eds.), Coastal and Marine Hazards, Risks and Disasters. London, GB: Elsevier, pp. 243–270.*

Source: Robert J. Nicholls, 'Adapting to Sea-Level Rise', in *Resilience: The Science of Adaptation to Climate Change*, ed. Zinta Zommers and Keith Alverson (Elsevier, 2018).

As this shows, some cities such as Tokyo have managed to adapt to sea level rise of up to 4m (40mm per year). Subsidence rates are extremely high in some places. In Jakarta, subsidence is at 100 millimetres per year, or 10 metres per century.²⁹²

Conversely, some places such as Helsinki have seen negative relative sea level rise due to glacial uplift.²⁹³

While the average coastal *area* experiences relative sea level rise of less than 3mm per year, the average coastal *resident* experiences a rise of around 9mm per year, due to subsidence. This is because coastal residents are concentrated in areas experiencing faster relative sea level rise.

²⁹² Hasanuddin Z. Abidin et al., 'Land Subsidence of Jakarta (Indonesia) and Its Relation with Urban Development', *Natural Hazards* 59, no. 3 (11 June 2011): 1753, <https://doi.org/10.1007/s11069-011-9866-9>.

²⁹³ Nicholls, 'Adapting to Sea-Level Rise', fig. 2.2.

Table 1 | Contribution of the climate and geologic components to relative sea-level change for length-weighted and population-weighted cases, respectively

Relative SLR component	Contribution to relative sea-level change			
	Coastal-length weighted		Coastal-population weighted	
	mm yr ⁻¹	%	mm yr ⁻¹	%
Climate-induced SLR (1993 to 2015)	3.2	122	3.8	39 to 49
GIA	-0.8	-32	-0.3	-3
Delta subsidence	0.1	4	1.6	16 to 21
City subsidence	0.1	3	2.7 to 4.8	35 to 49
Global-mean sum	2.6		7.8 to 9.9	

Average values are reported, except for cities and the global-mean sum where a low/high range is used to express the uncertainty in subsidence (Supplementary Table 2).

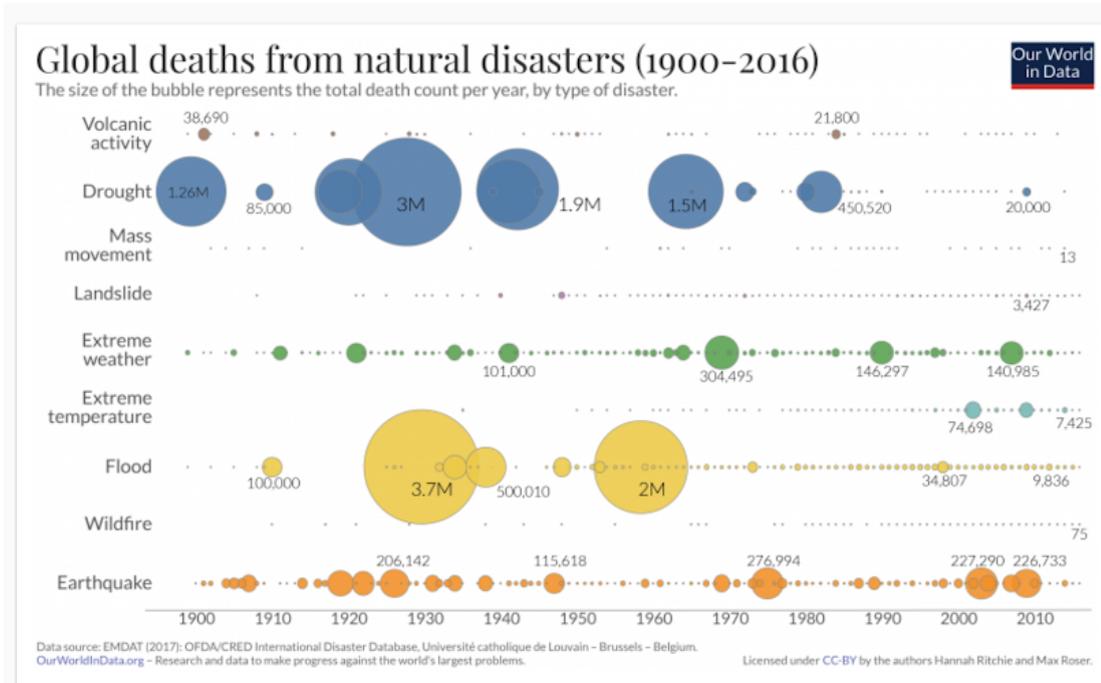
Source: Robert J. Nicholls et al., 'A Global Analysis of Subsidence, Relative Sea-Level Change and Coastal Flood Exposure', *Nature Climate Change*, 8 March 2021, 1–5, <https://doi.org/10.1038/s41558-021-00993-z>.

7.1.3. Floods

One of the major impacts of relative sea-level rise are more frequent and intense coastal floods, because higher mean sea-levels raise extreme sea-levels (tides, surges, waves), which then propagate inland causing losses and damages. While deaths from coastal floods declined dramatically over the course of the 20th century due to improved disaster risk reduction,²⁹⁴ economic losses continue to rise, mainly due to increased exposures and, locally, due to subsidence.²⁹⁵

²⁹⁴ Bouwer, L.M., Jonkman, S.N., 2018. Global mortality from storm surges is decreasing. *Environmental Research Letters* 13, 014008. <https://doi.org/10.1088/1748-9326/aa98a3>.

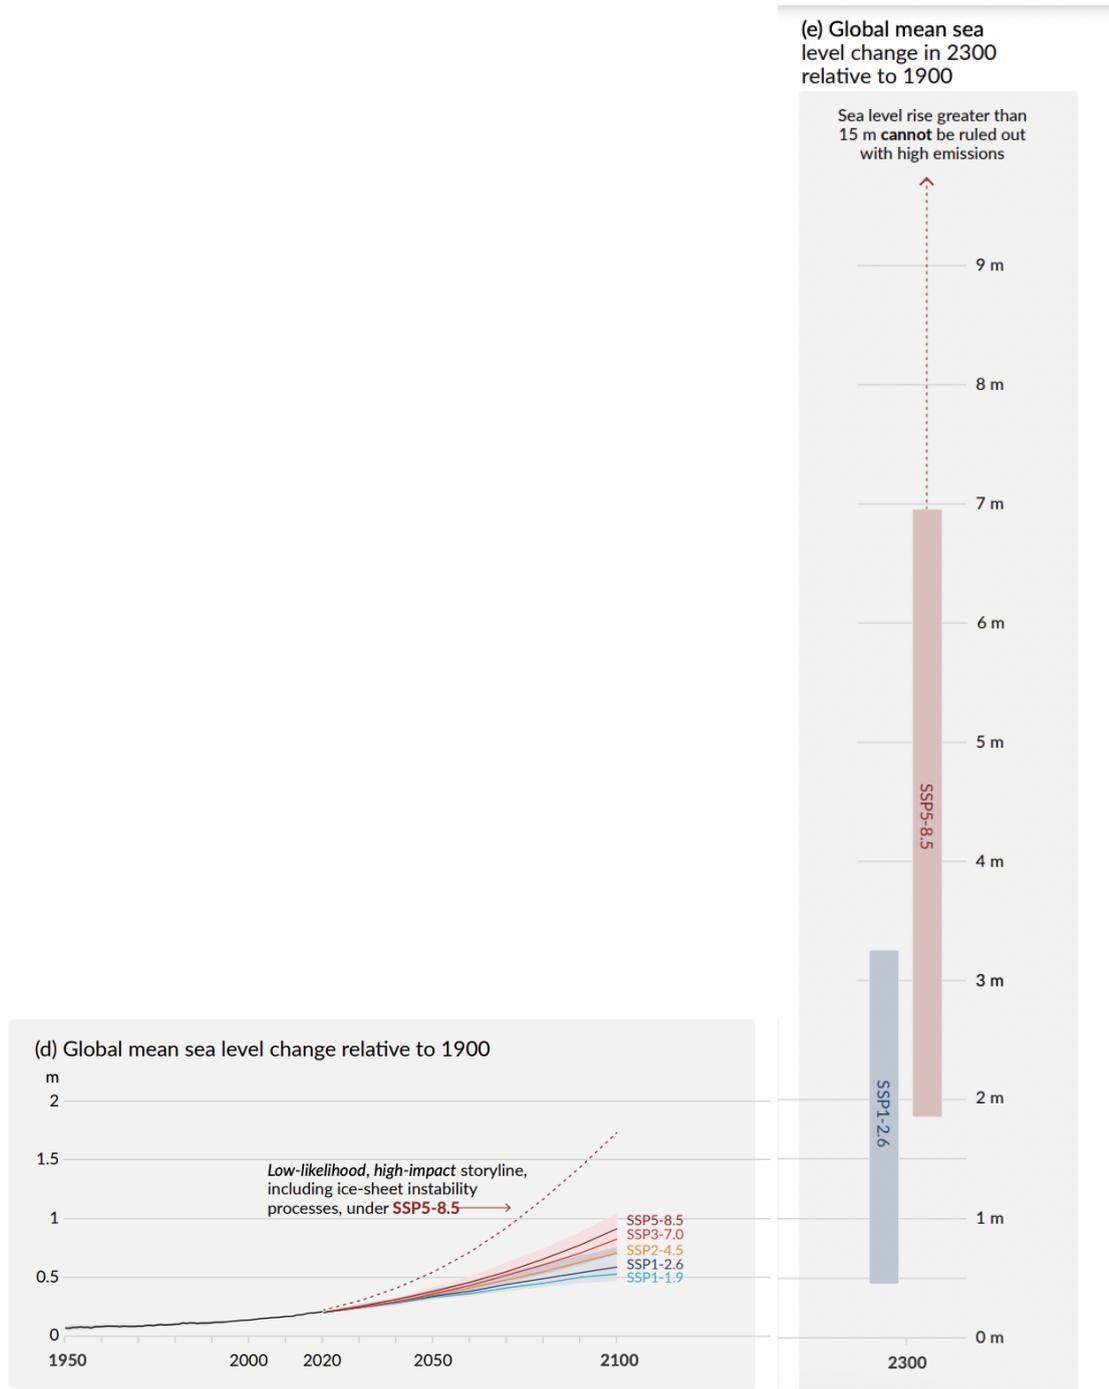
²⁹⁵ Kron, W., 2013. Coasts: the high-risk areas of the world. *Nat Hazards* 66, 1363–1382. <https://doi.org/10.1007/s11069-012-0215-4>.



7.2. How much will sea level rise due to global warming?

The potential impact of future sea level rise illustrates two important points: (1) we need to pay attention to low probability but high-impact risks, and (2) taking the long-term perspective can fundamentally change how we should assess global risks.

The figure below shows the 17-83% confidence range of sea level rise on different emissions scenarios to 2100 and to 2300:



Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021). SPM.8.

The IPCC does not say how likely the suggested 15 metre sea level rise might be on RCP8.5, but only that it “cannot be ruled out”. For this very high emissions scenario, deep uncertainty is a real concern and the IPCC has low confidence that the models capture all of the processes at play.²⁹⁶ Various studies have explored the implications of the rapid collapse

²⁹⁶ “By contrast, for SSP5-8.5, the SEJ and MICI projections exhibit 17th -83rd percentile ranges of 0.02-0.56 m and 0.19-0.53 m by 2100, consistent with one another but considerably broader than the likely contribution for medium confidence processes of 0.03 to 0.34 m. This lower level of agreement for higher emissions scenarios reflects the deep uncertainty in the AIS contribution to GMSL change under higher emissions scenarios (Box 9.4). This deep uncertainty grows after 2100: by 2150, under

of the West Antarctic Ice Sheet, causing 5m of sea level rise over 100 years (50mm per year).²⁹⁷

Over longer timescales, the effects on sea level rise will be even more pronounced. Sea levels will continue to rise in the millennia after emissions stop, as the ice sheets melt. The table below compares observed peak rates and magnitudes of relative sea level rise from past observations and for future projections under different scenarios:

Location	Period	Considered drivers of rel sea-level rise	Scenario	Peak rate of rel sea level rise	Magnitude of rel sea level rise
Past observations					
New Orleans	20th Century	Subsidence	N/A	25mm/year	2.5m ²⁹⁸
Tokyo	20th Century	Subsidence	N/A	40mm/year	4m ²⁹⁹
Jakarta	1982-2010	Subsidence	N/A	100mm/year	2.8m ³⁰⁰
Global coastal population	1993-2015	Subsidence + climate change	N/A	9mm/year	2m ³⁰¹
Future projections					
Global mean	To 2100		2.7°C	8mm/year	0.7m ³⁰²
Global mean	To 12,000 AD		2.7°C	10mm/year	10-20m ³⁰³

SSP 5-8.5, medium confidence processes likely lead to a -0.1 to 0.7 m AIS contribution, while SEJ and MICI-based projections indicate 0.0-1.1 m and 1.4-3.7 m, respectively” IPCC, *Climate Change 2021: The Physical Science Basis, Sixth Assessment Report* (UNFCCC, 2021), ch. 9, sec. 9.6.3.2.

²⁹⁷ For a review, see Robert J. Nicholls, Richard S. J. Tol, and Athanasios T. Vafeidis, ‘Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet: An Application of FUND’, *Climatic Change* 91, no. 1 (25 June 2008): 171, <https://doi.org/10.1007/s10584-008-9424-y>.

²⁹⁸ Robert J. Nicholls, ‘Adapting to Sea-Level Rise’, in *Resilience: The Science of Adaptation to Climate Change*, ed. Zinta Zommers and Keith Alverson (Elsevier, 2018).

²⁹⁹ Nicholls, ‘Adapting to Sea-Level Rise’.

³⁰⁰ Nicholls, ‘Adapting to Sea-Level Rise’.

³⁰¹ Robert J. Nicholls et al., ‘A Global Analysis of Subsidence, Relative Sea-Level Change and Coastal Flood Exposure’, *Nature Climate Change*, 8 March 2021, 1–5, <https://doi.org/10.1038/s41558-021-00993-z>.

³⁰² IPCC, *Climate Change 2021: The Physical Science Basis, Sixth Assessment Report* (UNFCCC, 2021), SPM.8.

³⁰³ Clark et al. (2016, fig. 4a) project that on a medium-low ‘emissions scenario’ close to RCP4.5, sea level would rise by 20 metres, while Van Breendam et al. (2020, Tab. 1) find that it would rise by 10 metres. Peter U. Clark et al., ‘Consequences of Twenty-First-Century Policy for Multi-Millennial Climate and Sea-Level Change’, *Nature Climate Change* (8 February 2016), <https://doi.org/10.1038/nclimate2923>; Jonas Van Breendam, Heiko Goelzer, and Philippe Huybrechts, ‘Semi-Equilibrated Global Sea-Level Change Projections for the next 10,000 Years’, *Earth System Dynamics* 11, no. 4 (6 November 2020): 953–76, <https://doi.org/10.5194/esd-11-953-2020>.

Global mean	To 2100		4.4°C	15mm/year	0.9m ³⁰⁴
Global mean	To 2120		Collapse of West Antarctic Ice Sheet	50mm/year	5m ³⁰⁵
Global mean	To 12,000 AD		All fossil fuels	10mm/year	20-30m ³⁰⁶

This shows how much of a difference taking a longtermist perspective can make: on *business as usual* - roughly, RCP4.5 - it is the difference between a 0.75 metre rise over 100 years and a >10 metre rise over 10,000 years.

7.3. What is the scope to adapt to relative sea level rise?

7.3.1. Adaptation options

The main adaptive responses are:³⁰⁷

- **Protection** - blocking the inland propagation of mean and extreme sea-levels by, e.g. building sea walls or dikes.
- **Advance/attack** - creation of new land by building into the sea.
- **Accommodation** - adapting settlements to higher water levels e.g. insurance, flood-proofing houses or putting them on stilts.
- **Retreat** - abandoning current settlements.

³⁰⁴ IPCC, Climate Change 2021: The Physical Science Basis, Sixth Assessment Report (UNFCCC, 2021). SPM.8.

³⁰⁵ Robert J. Nicholls, Richard S. J. Tol, and Athanasios T. Vafeidis, 'Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet: An Application of FUND', *Climatic Change* 91, no. 1 (25 June 2008): 171, <https://doi.org/10.1007/s10584-008-9424-y>.

³⁰⁶ Peter U. Clark et al., 'Consequences of Twenty-First-Century Policy for Multi-Millennial Climate and Sea-Level Change', *Nature Climate Change* (8 February 2016): fig. 2b, <https://doi.org/10.1038/nclimate2923>.

³⁰⁷ IPCC, 'Special Report on the Ocean and Cryosphere in a Changing Climate', sec. 4.1.4.

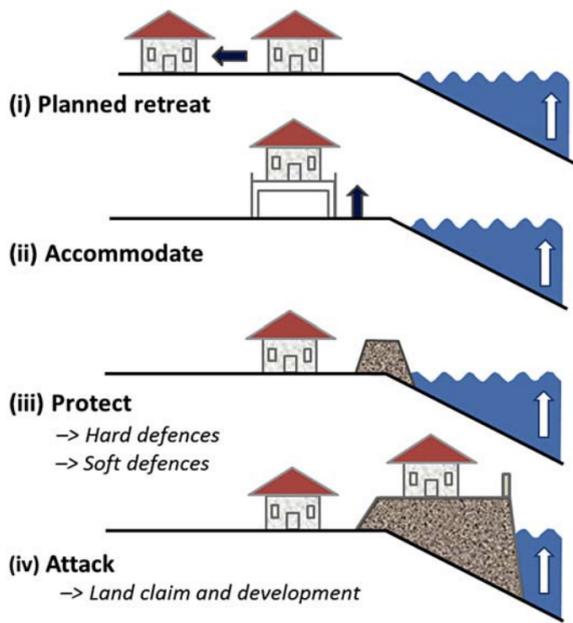


FIGURE 2.5 Generic adaptation approaches for sea-level rise, modified to include attack (via land claim and coastal development). Modified from Dronkers, J., Gilbert, J.T.E., Butler, L.W., Carey, J.J., Campbell, J., James, E., et al., 1990. *Strategies for adaption to sea level rise. Report of the Coastal Zone Management Subgroup, Response Strategies Working Group of the Intergovernmental Panel on Climate Change. The Hague, The Netherlands, Ministry of Transport, Public Works and Water Management.*

Source: Nicholls, 'Adapting to Sea-Level Rise'.

7.3.2. Future adaptation prospects

The scope to adapt to future relative sea level rise depends on the rate and magnitude of sea level rise.

On a population basis, future sea level rise mainly threatens Asian countries and island nations.

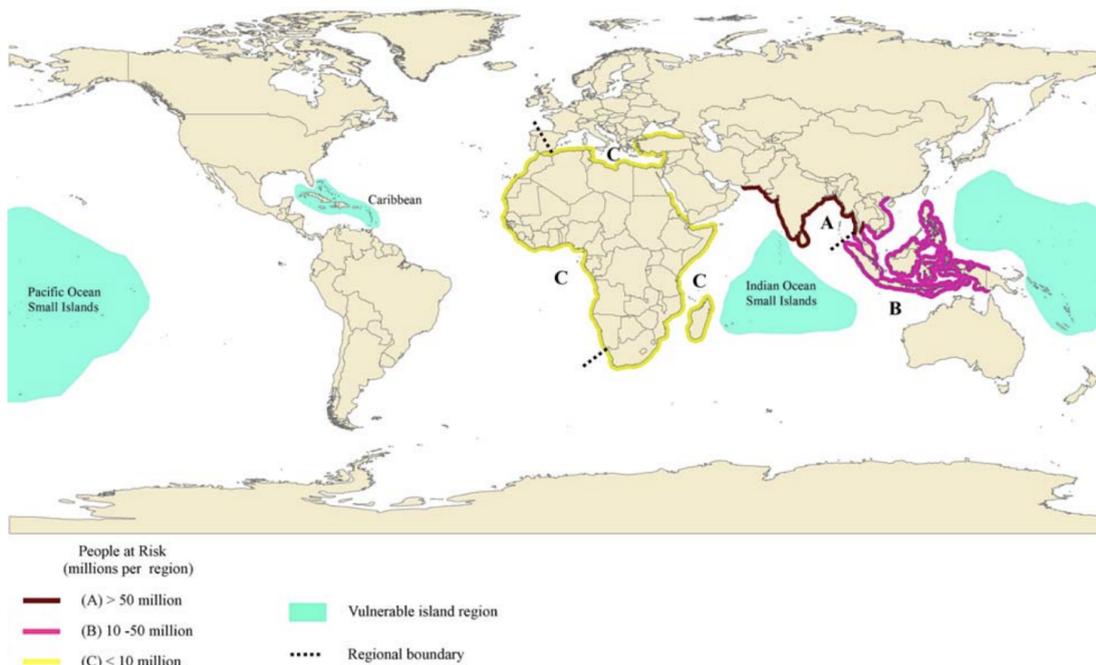


FIGURE 2.4 Regions most vulnerable to coastal flooding and sea-level rise. At highest risk are coastal zones with dense populations, low elevations, appreciable rates of subsidence, and/or inadequate adaptive capacity. From Nicholls, R.J., Cazenave, A., 2010. *Sea level rise and its impact on coastal zones. Science, 328, 1517–1520 (Nicholls and Cazenave, 2010).*

Source: Robert J. Nicholls, 'Adapting to Sea-Level Rise', in *Resilience: The Science of Adaptation to Climate Change*, ed. Zinta Zommers and Keith Alverson (Elsevier, 2018).

I will now discuss the prospects for adaptation in different scenarios.

Up to 2 metres over 80 years

The picture that emerges from the literature on adaptation to relative sea level rise is that there are few *technical* limits to sea level rise of up to 2 metres by 2100 (25mm per year on average).

Table 1 | The coastal and social characteristics and adaptation constraints to maintaining human settlements safe from twenty-first-century sea-level rise for cases considered in this Perspective

Case	Dominant coastal characteristics			Adaptation constraints				
	Coastal landform	World Bank income group in 2017	Human settlements	Mean population density (people km ⁻²)*	Technological limits	Economic barriers	Finance barriers	Social conflict barriers
Bangladesh	Delta	Lower middle income	Rural	^a 1,100	—	—	X	X
Catalonia	Beaches, deltas, cliffs	High income	Rural/urban	^a 900	—	—	—	X
Ho Chi Minh City	Delta	Lower middle income	Urban	^a 3,900	—	—	X	X
Maldives	Atoll islands	Higher middle income	Urban	^b 26,000	—	—	—	X
			Rural	^b 900	—	X	X	X
New York City	Estuary	High income	Urban	^a 11,000	—	—	X	X
Netherlands	Delta, beaches	High income	Rural/urban	^a 500	—	—	—	X

Dashes (—) and crosses (X) denote the absence and presence of adaptation constraints across the set of available adaptation options, respectively. *Mean population density values are based on the UN-adjusted GPWv4 year 2010 population density dataset³⁰⁸ and the Global Administrative Areas (GADM) dataset version 2.0 (<http://www.gadm.org/>). For New York City and Ho Chi Minh City mean population density was calculated for the entire administrative area of the city. For Catalonia, Bangladesh and the Netherlands mean population density was calculated for the LE CZ (low elevation coastal zone; areas ≤10 m and hydrologically connected to the ocean). For Bangladesh, the districts of Cox's Bazar, Bandarban, Chittagong, Ramgamati and Khagrachhari were excluded as they are outside the delta. For the definition of the LE CZ we used CGIAR-CSI SRTM v4.1 elevation data³⁰⁹. ^bPopulation counts for the Maldives are taken from the Maldives Population and Housing Census 2014 (<http://statistics.maldives.gov.mv/nbs/wp-content/uploads/2015/12/PP5.xls>). We define urban as population living in the city of Malé and rural as population living on other atolls. We use GADM version 2.0 to define the administrative boundaries of the Maldives and Malé. ^cEconomic and financial barriers may arise for maintaining beaches not used for tourism.

Source: Jochen Hinkel et al., 'The Ability of Societies to Adapt to Twenty-First-Century Sea-Level Rise', *Nature Climate Change* 8, no. 7 (July 2018): 570–78, <https://doi.org/10.1038/s41558-018-0176-z>.

Generally, rich countries will have the incentives and the resources to invest in sea level adaptation, and in particular coastal protection. The IPCC notes that:

“If governments undertook adaptation investments in all coasts (e.g., building protective dikes), then the study suggests... a population of less than half a million displaced under the 2.0 m sea level rise scenario.”³⁰⁸

Such protection measures are likely to be implemented given the long history of coastal protection (not least due to subsiding cities as discussed above), high benefit-cost ratios of coastal protection and the high cost of not protecting urban centres.³⁰⁹

³⁰⁸ IPCC, *Climate Change 2014: Impacts, Adaptation, and Vulnerability: Summary for Policymakers* (Cambridge University Press, 2014), chap. 12.

³⁰⁹ Lincke, D., Hinkel, J., 2018. Economically robust protection against 21st century sea-level rise. *Global Environmental Change* 51, 67–73. <https://doi.org/10.1016/j.gloenvcha.2018.05.003>

However, for poorer and very exposed countries, including island nations and some countries in Africa and Asia, adaptation could in some cases be too costly, so abandonment might be the chosen path.³¹⁰

“Effective protection requires investments on the order of tens to several hundreds of billions of USD yr⁻¹ globally (high confidence). While investments are generally cost efficient for densely populated and urban areas (high confidence), rural and poorer areas will be challenged to afford such investments with relative annual costs for some small island states amounting to several percent of GDP (high confidence). Even with well-designed hard protection, the risk of possibly disastrous consequences in the event of failure of defences remains.”³¹¹

As I argued in Chapter 4, it is inappropriate to discount future costs on the basis that future people will be better off because there is a good chance that many people in Asia and Africa will not actually be much better off.

Moreover, as shown in the table above, there are social and political barriers which could lead to suboptimal adaptation even in rich countries.

Relative sea level rise of 5 metres over 100 years

If the Western Antarctic Ice Sheet collapses, there would be 5 metres of sea level rise over the course of 100 years, which is around 50mm per year. Would we be able to adapt to such an extreme event?

Past experience with subsidence and living below sea level

Some cities and regions have adapted to multimetre relative sea level rise caused by subsidence, which provides some insights on what could happen in the future under 5 metre of sea-level rise. For example, Tokyo adapted to 4m of relative sea level rise over the 20th century (40mm per year).

So far, retreat has been an unpopular option and has mainly been reserved for small communities or carried out to create new wetland habitat.³¹² Rather, *successful* protection has been the dominant response in almost all populated places suffering from subsidence, which usually also enjoy high economic and population growth.³¹³ One exception to this is

³¹⁰ “Vulnerability to sea-level rise is not uniform and small islands, Africa and south, southeast and east Asia are recognized as the most vulnerable regions [11]. This reflects their high and growing exposure and low adaptive capacity. These regions are the areas where protection is most likely to not occur or fail, and they collectively contain a significant proportion of potential environmental refugees, especially the Asian regions (figure 3). Many of the people in Asia live in deltas, which are extensive and often subsiding coastal lowlands, amplifying global changes and making them more challenging environments for adaptation [47,48,82]. Small islands have relatively small population and given that implementing protection could also present significant problems, forced abandonment seems a feasible outcome for small changes in sea level” Robert J. Nicholls et al., ‘Sea-Level Rise and Its Possible Impacts given a “beyond 4°C World” in the Twenty-First Century’, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369, no. 1934 (13 January 2011): 161–81, <https://doi.org/10.1098/rsta.2010.0291>.

³¹¹ IPCC, ‘Special Report on the Ocean and Cryosphere in a Changing Climate’, 2019, 56, <https://www.ipcc.ch/srocc/>.

³¹² IPCC, ‘Special Report on the Ocean and Cryosphere in a Changing Climate’, 55.

³¹³ Nicholls, ‘Adapting to Sea-Level Rise’.

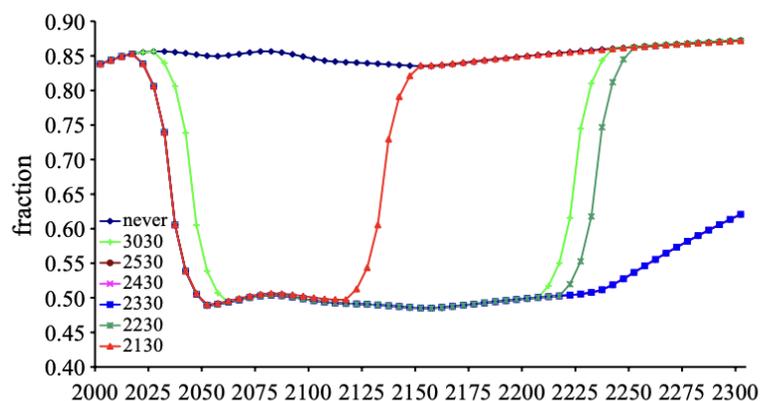
New Orleans, which experienced relative sea level rise of 2-3m over the 20th century due to subsidence (20-30mm per year).³¹⁴ Despite having the technical and economic capacity to protect New Orleans, the defences were breached during Katrina and the population levels have not returned to pre-Katrina levels.³¹⁵

Currently, “at least 20 million people accept the risk of living up to several metres below normal high tides in countries such as Belgium, Canada, China, Germany, Italy, Japan, the Netherlands, Poland, Thailand, the United Kingdom and the United States”.³¹⁶

Future projections of adaptation

The Atlantis Project explored the scope to adapt to 5-6 metres of sea level rise over the course of 100 years (50mm per year). This could be caused by the collapse of the West Antarctic Ice Sheet. Given certain assumptions about the costs of adaptation, “the length of the world’s coast that is protected declines from about 85% to about 50% of the exposed and populated coastline, reflecting that protection becomes too expensive in many areas”.³¹⁷ However, because population is unevenly distributed along the coasts, this would still protect >95% of the coastal population.³¹⁸

Fig. 4 The fraction of coastal protection of exposed and populated coasts at the global scale as a function of time and WAIS scenario



Source: Nicholls, Tol, and Vafeidis, ‘Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet’.

At a rate of only 5mm per year, on this model, it is still true that much of the world’s coastline would be abandoned (this is shown by the green line).

Regional studies suggest that the response would be a mix of protection, accommodation and retreat even in rich countries.

³¹⁴ Robert J. Nicholls et al., ‘A Global Analysis of Subsidence, Relative Sea-Level Change and Coastal Flood Exposure’, *Nature Climate Change*, 8 March 2021, 1–5, <https://doi.org/10.1038/s41558-021-00993-z>.

³¹⁵ The population was 390,000 in 2019 compared to 500,000 before Katrina. Nicholls, ‘Adapting to Sea-Level Rise’, 19.

³¹⁶ Hinkel et al., ‘The Ability of Societies to Adapt to Twenty-First-Century Sea-Level Rise’, 575.

³¹⁷ Robert J. Nicholls, Richard S. J. Tol, and Athanasios T. Vafeidis, ‘Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet: An Application of FUND’, *Climatic Change* 91, no. 1 (25 June 2008): 171, <https://doi.org/10.1007/s10584-008-9424-y>.

³¹⁸ Nicholls, Tol, and Vafeidis, 187.

Table 3 Summary of the response of three detailed case studies to the Atlantis sea-level rise scenario (5-m rise from 2030 to 2130)

Case Study	Impacts	Response	Source
Netherlands	>10 million people threatened together with one of the world's largest economies	Abandon the northwest and southwest of the Netherlands, but possibly protect the Ranstad (Amsterdam to Rotterdam area). Likelihood of intense political conflict and very large response costs in proportion to GDP	Olsthoorn et al. (2008)
Thames estuary	Two million people threatened with a rapidly increasing flood risk without a response, even allowing for expected upgrades. Also much of London's financial sector including Canary Wharf.	Indecision may lead to forced abandonment, but there are adaptation options—especially a new downstream barrier	Lonsdale et al. (2008) Dawson et al. (2005)
Rhone delta	Compared to the other case study sites, human impacts are minimal, but significant natural values are threatened	After an initial 'wait and see', abandon the delta	Poumadere et al. (2008)

Source: Nicholls, Tol, and Vafeidis.

This is driven not by economic cost or technical feasibility but rather by political and social barriers which preclude the implementation of optimal policy.³¹⁹

One study suggests that London could mitigate the risk of 8m of sea level rise by moving the Thames Barrier to Canvey.³²⁰ There is research showing that the Netherlands could adapt to up to 5 metres of sea level rise using current engineering technology.³²¹

7.3.3. Effects on 21st Century sea level rise on land loss and migration

One important way that sea level rise might drive domestic and international political tension is by causing migration from coastal areas. Lincke and Hinkel (2021) explore cumulative land loss and cumulative displacement over the 21st Century on different warming scenarios and on different Shared Socioeconomic Pathways. They consider five sea level rise scenarios ranging from 33–170 cm in 2100 (RCP2.6 low ice melt, RCP2.6 high ice melt, RCP8.5 low ice melt, RCP8.5 high ice melt and high-end).

³¹⁹ "Although the Netherlands and the UK have the technological and economic wherewithal to adapt to extreme sea-level rise, the case studies suggest that this necessary condition is not a sufficient one, as assumed by the model." Nicholls, Tol, and Vafeidis.

³²⁰ Jim W. Hall, Hamish Harvey, and Lucy J. Manning, 'Adaptation Thresholds and Pathways for Tidal Flood Risk Management in London', *Climate Risk Management* 24 (1 January 2019): 42–58, <https://doi.org/10.1016/j.crm.2019.04.001>.

³²¹ "Research also showed that with an investment of around €80 billion, it may be possible to preserve territorial integrity of the Netherlands even under 5 m of SLR, using current engineering technology" Hinkel et al., 'The Ability of Societies to Adapt to Twenty-First-Century Sea-Level Rise', 574.

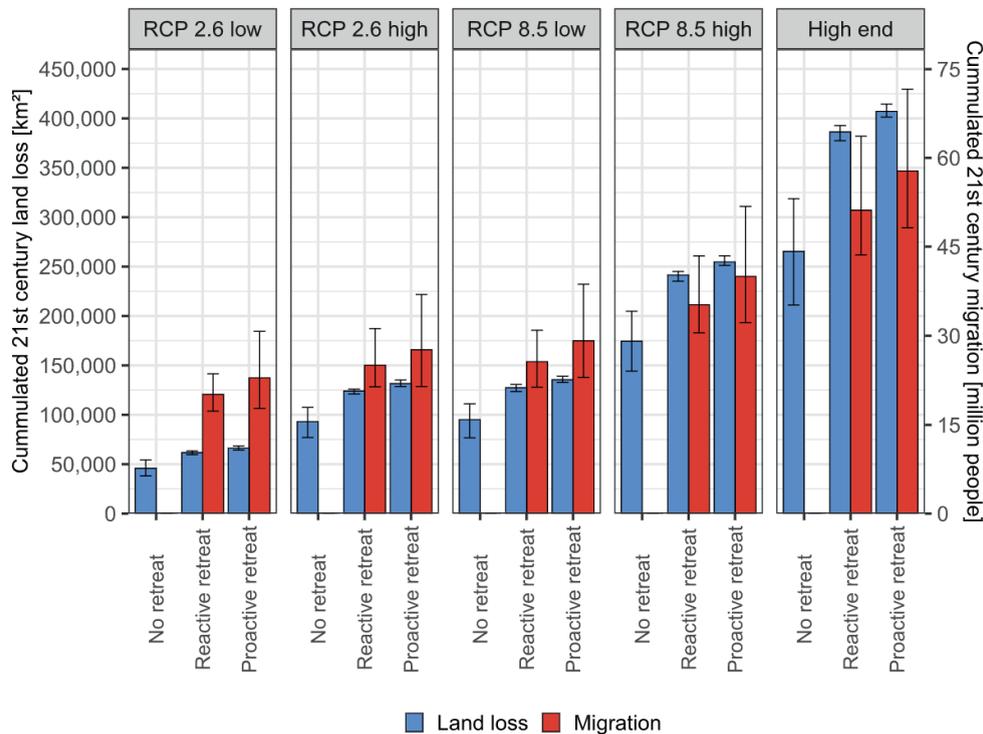


Figure 1. Global land loss and migration cumulative over the 21st century under different retreat assumptions and the five sea-level rise (SLR) scenarios used in this study. The bars show averages over all shared socioeconomic pathway (SSP) scenarios and discount rates, the error bars show the associated uncertainty range.

Source: Daniel Lincke and Jochen Hinkel, ‘Coastal Migration Due to 21st Century Sea-Level Rise’, *Earth’s Future* 9, no. 5 (2021): e2020EF001965, <https://doi.org/10.1029/2020EF001965>.

For around 1 metre of sea level rise (RCP8.5 high scenario),³²² 30-45 million people will migrate cumulatively over the 21st Century, depending on the socioeconomic scenario. This is the cumulative effect over the whole 21st Century. So, this works out at around 375,000 to 600,000 people per year. If sea level rise is limited to 30cm (RCP2.6 low), then around 20 million cumulatively would be displaced. So, 1 metre of sea level rise compared to 30cm increases the cumulative number of displaced people by up to 25 million, or around an extra 310,000 per year.

Migration will mainly be concentrated in Asia:

³²² Lincke and Hinkel, ‘Coastal Migration Due to 21st Century Sea-Level Rise’, Table S2.

Table 2
Cumulative Land Loss and Forced Migration During 21st Century for the Ten Most Affected countries

Land loss (km ²)			
Rank	Country	Median	Min-max range
1	Russia	28,600	13,600–56,600
2	Canada	23,400	8,000–43,900
3	USA	15,200	8,100–39,000
4	India	8,300	3,800–20,600
5	Brazil	7,200	2,500–25,700
6	Greenland	6,400	0–24,600
7	Australia	6,200	3,200–24,800
8	Viet Nam	4,900	200–16,200
9	Mexico	3,800	2,300–13,100
10	Indonesia	3,500	700–21,400
Migration (10 ⁶ people)			
1	Viet Nam	5.7	1.2–12.0
2	India	5.4	3.0–14.5
3	Indonesia	1.4	0.8–4.5
4	Bangladesh	1.2	0.2–8.0
5	Philippines	1.0	0.6–2.5
6	Brazil	0.9	0.5–1.8
7	Myanmar	0.9	0.4–3.0
8	China	0.8	0.2–3.3
9	South Korea	0.7	0.5–1.3
10	Japan	0.6	0.2–2.3

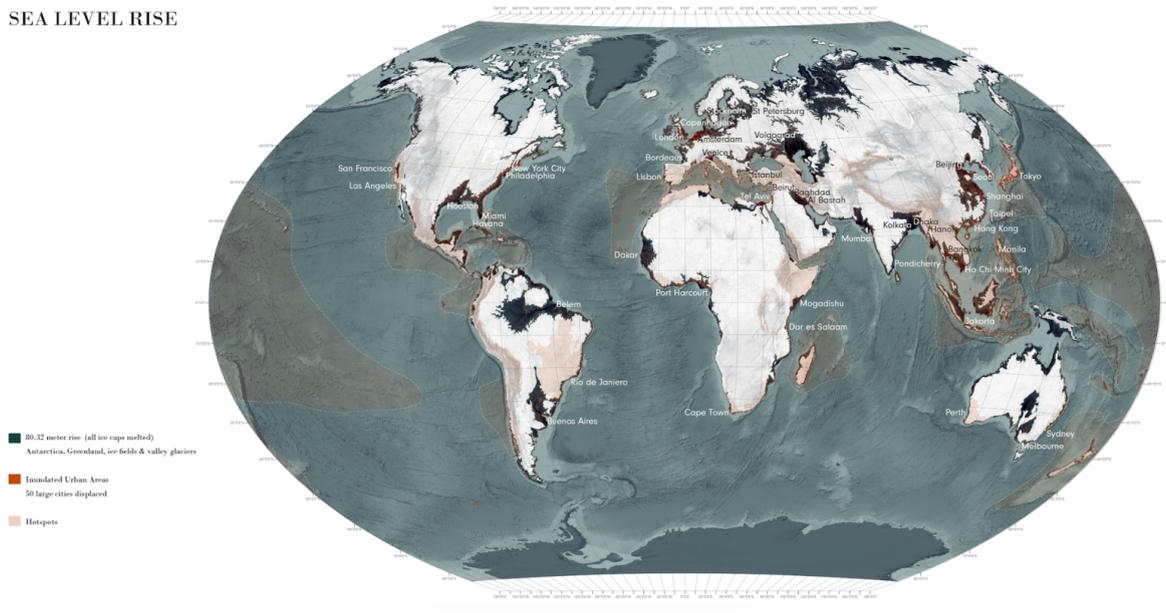
Note: The median column shows the median value over all scenarios, the min-max range column the minimum and maximum over all scenarios. Countries are selected and ordered according to the median value.

Source: Daniel Lincke and Jochen Hinkel, 'Coastal Migration Due to 21st Century Sea-Level Rise', *Earth's Future* 9, no. 5 (2021): e2020EF001965, <https://doi.org/10.1029/2020EF001965>.

As an upper bound on the effects of sea level rise, we can consider what would happen if all of the ice caps were to melt.³²³ The US Geological Survey [estimates](#) that this would cause sea levels to rise by 80 metres. The effects this would have are shown on the map below:

³²³ I owe this point to Benjamin Hilton's [review of climate change](#) for 80,000 Hours.

SEA LEVEL RISE



As the [Atlas for the End of the World](#) says, in this scenario “vast new coastlines and inland seas will be created and 50 of the world’s major cities would become architectural reefs”. But still, the vast majority of the world would be above water.

7.3.4. Economic effects of sea level rise

I discuss estimates of the economic costs of sea level rise in Chapter 10.

7.4. Overall verdict on sea level rise

Sea level rise illustrates two important insights that have been stressed by longtermist and effective altruist researchers. Firstly, we need to pay attention to low probability, high-impact events because these may account for most of the expected costs of sea level rise. While on RCP8.5, the most likely level of sea level rise in 2300 is around 4 metres, we cannot rule out a rise of 15 metres.

Secondly, if we ignore all impacts beyond 2100 we would in effect ignore some truly huge changes that future generations will have to deal with. On *business as usual*, by 2100, there would be around 75 cm of sea level rise. But in 10,000 years’ time, sea level would be more than 10 metres higher. Millennia into the future, the world will look very different.

Historical experience and modelling studies suggest that most coastal regions will successfully adapt to up to 2 metres of sea level rise by 2100. The technical barriers to successful adaptation seem low. Rich and densely populated areas are very likely to invest in the requisite adaptation measures given the large net benefits, though political and social factors might stand in the way of adaptation in some cases. Poorer, particularly exposed (i.e. small island states) and rural areas, however, will likely not be able to afford such protection and will thus be confronted with the need to migrate away from the coast.

Studies also suggest that adaptation to extreme scenarios, such as 5 metres of sea level rise due to the collapse of the West Antarctic Ice Sheet, is technically and economically feasible, though successful adaptation would be much harder.

In any case, two things must be emphasised.³²⁴ First, even if societies by and large manage to adapt to several metres of sea-level rise during the 21st century, such a scenario is likely to go along with massive human suffering among those who lose their homes and livelihoods due to retreat, or those being affected by coastal disasters in the case that coastal protection measures fail. Second, it must be emphasised that if sea-levels have already risen by 2 or more metres at the end of the century, sea-level rise will progress beyond 2100 at very high rates and will eventually also threaten those places that had managed to protect against 21st century sea-level rise.

³²⁴ Thanks to Jochen Hinkel for raising this point.

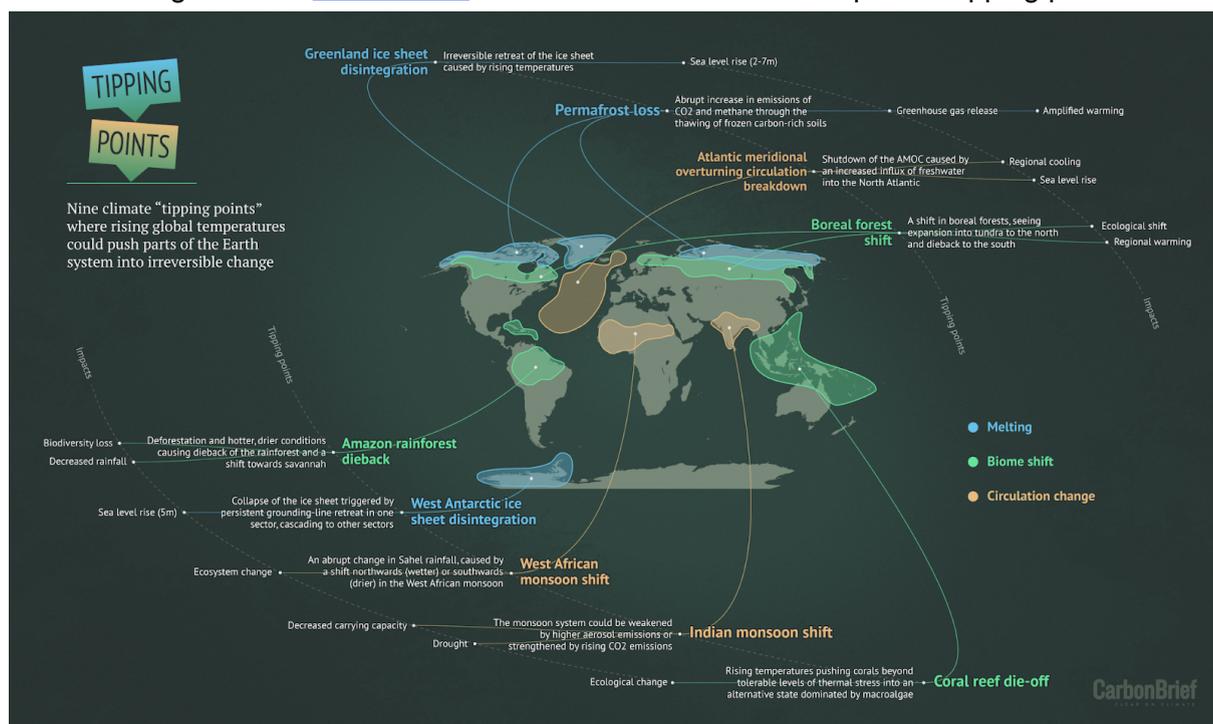
8. Tipping points

For those concerned about the long-term future and global catastrophic risk, potential tipping points or non-linearities in the climate system are especially important.

Non-linearities can be gradual or sudden. For instance, climate sensitivity might increase super-linearly in a relatively modest and gradual way, or there might be a dramatic step change at higher levels of warming or emissions. Sudden and dramatic non-linearities pose the most severe global catastrophic risks. Tipping points are abrupt changes in a system that are irreversible over a short timescale, such as a few decades. The impacts of tipping points need not be abrupt. For example, it might be that passing a certain level of warming will cause the Greenland ice sheet to break up over thousands of years, and that this effect cannot be reversed, except on millennial timescales. But the impact of this change might still be felt only over many thousands of years as the ice sheets slowly melt.

The tipping points with the greatest human impact will tend to be those that have most of their effects over the course of years to decades because we would have less time to adapt. For instance, rising sea levels over thousands of years would be bad, but give us lots of time to adapt.

This nice diagram from [CarbonBrief](#) outlines some of the most important tipping points.



8.1. Permafrost carbon release

[Permafrost](#) is ground that has been frozen for at least two consecutive years. Its thickness ranges from less than one metre to more than a kilometre. Typically, it sits beneath an “active layer” that thaws and refreezes every year. When temperatures rise, the permafrost may start to thaw. Permafrost thaw is one of the most frequently discussed potential tipping points.

About 1 trillion tonnes of carbon is stored in permafrost.³²⁵ The IPCC's Sixth Assessment Report estimates that for each 1°C of warming, permafrost emissions will increase by 18 GtC, with a 5% to 95% range of 3 to 42 GtC.³²⁶ For reference;

- Cumulative emissions from fossil fuel and industry since the Industrial Revolution = 464 GtC.
- Global emissions from fossil fuel and industry in 2019 = 10 GtC.
- Cumulative emissions on RCP4.5 (2019-2100) = 850 GtC.³²⁷

On RCP4.5, temperatures would increase by a further 1.5°C relative to today. This would increase permafrost emissions by 5 GtC to 60 GtC (5% to 95% range), which is important but small relative to cumulative anthropogenic emissions. An additional 60 GtC would add around 0.1°C to global average temperatures.³²⁸

The IPCC estimates that up to half of the permafrost carbon could be released abruptly, and the rest gradually.³²⁹

Beyond 2100, permafrost emissions would increase further. Different models produce different estimates of permafrost emissions on different anthropogenic emissions scenarios. Some of these are shown below:³³⁰

- RCP2.6 = 20-40 GtC.
- RCP4.5 = 17 GtC (range: minus 14 GtC to 54 GtC)
- RCP8.5 = 314 GtC (81 to 642 GtC)

So, on RCP4.5, permafrost emissions are small relative to anthropogenic emissions. On RCP8.5, at the upper end of the model range for 2100-2300, more than 600 GtC would be released, enough to warm the planet by about a degree. Still, on the extension of RCP8.5 to

³²⁵ “The new northern permafrost zone carbon inventory reports the surface permafrost carbon pool (0–3 m) to be 1,035 ±150 Pg carbon (mean ±95% confidence interval, CI).” E. a. G. Schuur et al., ‘Climate Change and the Permafrost Carbon Feedback’, *Nature* 520, no. 7546 (April 2015): 171–79, <https://doi.org/10.1038/nature14338>.

³²⁶ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box 5.1.

³²⁷ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, SPM. 8.

³²⁸ This is using the IPCC's TCRE estimate of 1.65C of warming per 1,000GtC.

³²⁹ “Abrupt thaw processes can contribute up to half of the total net greenhouse gas release from permafrost loss, the rest attributed to gradual thaw (Schneider von Deimling et al., 2015; Turetsky et al., 2020).” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box 5.1.

³³⁰ “Beyond 2100, models suggest that the magnitude of the permafrost carbon feedback strengthens considerably over the period 2100–2300 under a high-emissions scenario (Schneider von Deimling et al., 2015; McGuire et al., 2018). Schneider von Deimling et al., (2015) estimated that thawing permafrost could release 20–40 PgC of CO₂ in the period from 2100 to 2300 under a RCP2.6 scenario, and 115–172 PgC of CO₂ under a RCP8.5 scenario. The multi-model ensemble in (McGuire et al., 2018) project a much wider range of permafrost soil carbon losses of 81–642 PgC (mean 314 PgC) for an RCP8.5 scenario from 2100 to 2300, and of a gain of 14 PgC to a loss of 54 PgC (mean loss of 17 PgC) for an RCP4.5 scenario over the same period” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box 5.1.

2300, anthropogenic CO₂ emissions would be 5,000 GtC, so fossil fuel emissions remain much larger than potential permafrost emissions. It is important to bear in mind that, as discussed in Chapter 2, higher end recoverable fossil fuel estimates are 3,000 GtC, so this scenario may not be feasible.

The evidence from the paleoclimate also suggests that warming of around 4°C is unlikely to trigger the abrupt release of huge amounts of carbon from permafrost. During the Pliocene, higher latitudes were upwards of 10°C warmer than today,³³¹ but there is no sign of a huge carbon release over the course of years to decades.

In summary, the thawing of permafrost looks set to release additional carbon into the atmosphere, which will have important effects, but the effect is small relative to fossil fuel emissions.

8.2. Methane clathrates

[Methane clathrates](#) or hydrates are forms of ice that contain large amounts of methane. They form at low temperatures and high pressures in continental margin marine sediments or within and beneath permafrost. The total global clathrate reservoir is estimated to contain 1500–2000 GtC,³³² which, as we saw in Chapter 1, is close to many estimates of the carbon contained in recoverable fossil fuels. The release of vast amounts of methane from clathrates is often brought up in discussion of climate disaster scenarios.³³³ For instance, Whiteman et al (2013) stated:

“As the amount of Arctic sea ice declines at an unprecedented rate, the thawing of offshore permafrost releases methane. A 50-gigatonne (Gt) reservoir of methane, stored in the form of hydrates, exists on the East Siberian Arctic Shelf. It is likely to be emitted as the seabed warms, either steadily over 50 years or suddenly.”³³⁴

³³¹ See Chapter 3.

³³² “The total global clathrate reservoir is estimated to contain 1500–2000 PgC” IPCC, *Climate Change 2021: The Physical Science Basis*, sec. 5.4.9.1.3.

³³³ See for example “A more remote (but even more vivid) possibility, which in principle should also be included, is heat-induced releases of the even vaster offshore deposits of methane trapped in the form of clathrates.6 There is a very small and unknown (but decidedly nonzero) probability over the long run of having destabilized methane from these offshore clathrate deposits seep into the atmosphere if the temperature of the waters bathing the continental shelves increases just slightly. The amount of methane involved is huge, although it is not precisely known. Most estimates place the carbon-equivalent content of methane hydrate deposits at about the same order of magnitude as all other fossil fuels combined. Over the long run, a methane outgassing–amplifier process could potentially precipitate a disastrous strong positive feedback warming. If it occurred at all, such an event would likely take centuries to materialize because the presumed initiator would be the slow-acting gradual warming of ocean waters at the depths of the continental shelves. Thus, while it is a low-probability event that might only transpire centuries from now (if at all), the possibility of a climate meltdown is not just the outcome of a mathematical theory but has a real physical basis.” Martin L. Weitzman, “Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change,” *Review of Environmental Economics and Policy* 5, no. 2 (July 1, 2011): 275–92, <https://doi.org/10.1093/reep/rer006>.

³³⁴ Gail Whiteman, Chris Hope, and Peter Wadhams, “Climate Science: Vast Costs of Arctic Change,” *Nature*, July 24, 2013, <https://doi.org/10.1038/499401a>.

Methane stays in the atmosphere for around 10 years, eventually oxidising to carbon dioxide. [This](#) University of Chicago climate model allows you to test the effect of a given amount of methane on warming. It suggests that adding 50 billion tonnes of methane in a single slug will lead to warming of around 2°C over ten years, which would be very damaging.

The evidence and expert opinion suggests that an abrupt and massive release of methane is very unlikely.

8.2.1. Methane emissions so far

There is mixed evidence of increasing overall methane emissions from the permafrost region so far. Atmospheric measurements show no detectable trends in methane emissions from the permafrost regions over the past 30 years, though the IPCC has high confidence that observations understate methane emissions.³³⁵

8.2.2. Most methane would not reach the atmosphere

If methane hydrate does melt, most of it would not reach the atmosphere. As Carolyn Ruppel, chief scientist for the US Geological Survey's Gas Hydrates Project,

“If the methane released during gas hydrate degradation reaches the ocean, it would mostly be consumed by bacteria in the water column and not reach the atmosphere. In permafrost areas, degrading gas hydrate is usually deeply buried, so permafrost thaw is the more important contributor to greenhouse gas emissions.”³³⁶

8.2.3. The paleoclimate record

One way to test how much methane might be released from clathrates in the future is by looking at past episodes of warming in which there was Arctic warming or the Arctic was much warmer than today.

³³⁵ “Atmospheric measurements and inversions performed at the global and regional scales do not show any detectable trends in annual mean CH₄ emissions from the permafrost region over the past 30 years (Jackson et al., 2020; Saunio et al., 2020; Bruhwiler et al., 2021), consistent with atmospheric measurements in Alaska that showed no significant annual trends, despite significant increase in air temperature (Sweeney et al., 2016). Atmospheric inversions and biospheric models do not show any clear trends in CH₄ emissions for wetland regions of the high latitudes during the period 2000–2016 (Patra et al., 2016; Poulter et al., 2017; Jackson et al., 2020; Saunio et al., 2020). Large uncertainties on wetland extent and limited data constraints place low confidence in these modelling approaches.

The SROCC also assessed with high confidence that CH₄ fluxes have been under-observed due to their high variability at multiple scales in both space and time, and that there is a persistent mismatch between top-down and bottom-up methane budgets, with emissions calculated by upscaling ground observations typically higher than emissions inferred from large-scale atmospheric observations (Thornton et al., 2016a; Saunio et al., 2020).” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box 5.1.

³³⁶ Ruppel Carolyn D. and Kessler John D., “The Interaction of Climate Change and Methane Hydrates,” *Reviews of Geophysics* 55, no. 1 (February 8, 2017): 126–68, <https://doi.org/10.1002/2016RG000534>.

One relevant analogue is the warming from the Pleistocene into the Holocene, when temperatures increased by around 4°C. The IPCC concludes that

“In conclusion, several independent lines of evidence indicate that permafrost thaw did not release vast quantities of fossil CH₄ associated with the transient warming events of the [Last Deglacial Transition]. This suggests that large emissions of CH₄ from old carbon sources will not occur in response to future warming (medium confidence).”³³⁷

If there were going to be a large methane input from melting clathrates, that would also likely have happened during the Last Interglacial, when global average temperatures were 1°C above pre-industrial levels.³³⁸ In the Arctic, temperatures were around 1-2°C higher than today,³³⁹ and there is some evidence that the Arctic was perennially ice-free.³⁴⁰ Despite that, there is no evidence of a massive release of methane from clathrates in this period.

There is disagreement about how much of a role, if any, the melting of methane hydrate played in the Paleocene-Eocene Thermal Maximum.

“There is low to medium confidence in evaluations of the total amount of carbon released during the PETM, as proxy data constrained estimates vary from around 3000 to more than 7000 PgC, with methane hydrates, volcanic emissions, terrestrial and/or marine organic carbon, or some combination thereof, as the probable sources of carbon (Zeebe et al., 2009; Cui et al., 2011; Gutjahr et al., 2017; Elling et al., 2019;

³³⁷ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box 5.1.

³³⁸ “In summary, GMST during the warmest millennia of the 38 LIG (within the interval of around 129–125 ka) is estimated to have reached 0.5°C–1.5°C higher values than 39 the 1850–1990 reference period (medium confidence)” IPCC, *Sixth Assessment Report: Working Group I The Physical Science Basis*, sec. 2.3.1.1.1.

³³⁹ In the Last Interglacial, temperatures were around 4-5C above pre-industrial levels. Today, temperatures are around 3C warmer in the Arctic than pre-industrial levels. “Stronger LIG summertime insolation at high northern latitudes drove Arctic land summer temperatures 4–5°C higher than in the pre-industrial era.” Maria-Vittoria Guarino et al., ‘Sea-Ice-Free Arctic during the Last Interglacial Supports Fast Future Loss’, *Nature Climate Change* 10, no. 10 (October 2020): 928–32, <https://doi.org/10.1038/s41558-020-0865-2>. On recent Arctic warming, see the [Arctic Report Card](#).

³⁴⁰ “While knowledge of past Arctic temperatures is robust, thanks to the available observations^{2,10}, the interpretation of Arctic sea-ice changes during the LIG has previously been afflicted by uncertainty^{8,10,12,13}. Water-isotope measurements from ice cores have been interpreted to suggest that, alongside the Arctic warming, there was a reduction in the mean annual sea-ice area⁸. Microfauna in LIG marine sediments recovered from boreholes on the Beaufort Sea Shelf have been interpreted as implying a lack of perennial Arctic sea-ice cover¹⁴, as have planktonic foraminifera recovered from some Arctic marine cores^{15,16}. Similarly, ostracodes on the Lomonosov and Mendeleev Ridges and Morris Jesup Rise have been interpreted as indicative of minimum sea-ice coverage during peak LIG warmth¹⁷. However, measurements of the recently developed sea-ice proxy IP25 (a carbon-25 highly branched isoprenoid lipid), when combined with terrestrial and open-water phytoplankton biomarkers, have been interpreted as evidence of perennial LIG ice cover in the central part of the Arctic Ocean¹³. While aspects of this particular application of IP25 are debated¹⁸, this result (see also Methods), along with the fact that no coupled climate models have simulated an ice-free Arctic during the LIG (refs. 10,11,13,19), has meant that the research community has spent considerable time debating whether or not summer sea ice disappeared during this important past warm period^{8,12,13,19}.” Maria-Vittoria Guarino et al., ‘Sea-Ice-Free Arctic during the Last Interglacial Supports Fast Future Loss’, *Nature Climate Change* 10, no. 10 (October 2020): 928–32, <https://doi.org/10.1038/s41558-020-0865-2>

S.M. Jones et al., 2019; Haynes and Hönisch, 2020). Methane emissions related to hydrate/permafrost thawing and fossil carbon oxidation may have acted as positive feedbacks (Lunt et al., 2011; Armstrong McKay and Lenton, 2018; Lyons et al., 2019), as the inferred increase in atmospheric CO₂ can only account for approximately half of the reported warming (Zeebe et al., 2009).³⁴¹

The PETM occurred against a background that was 12°C warmer than today. If methane clathrates did play a role, they would have been destabilised by the release of more than 1.5 trillion tonnes of carbon, mainly in the form of carbon dioxide, on top of a much warmer baseline.³⁴²

Overall, there is little indication from the paleoclimate record that warming of 4°C relative to pre-industrial will cause an abrupt and massive release of methane from clathrates, as posited by Whiteman et al. The paleoclimatic evidence suggests that methane clathrates may become more of a concern if warming passes 10°C above pre-industrial levels and/or carbon emissions are huge.

The anti-climate sceptic website Skeptical Science provides an [accessible overview](#) of some of these arguments.

8.2.4. Expert estimates of emissions from methane hydrates

The view expressed by Whiteman et al (2013) is at odds with the consensus in the literature and has been the subject of significant criticism. There have also been two responses in *Nature* criticising the paper.³⁴³ In a review article, Schuur et al comment that

“A large pulse release of permafrost carbon on this timescale could cause climate change that would incur catastrophic costs to society, but there is little evidence from either current observations or model projections to support such a large and rapid pulse.”³⁴⁴

The IPCC’ Sixth Assessment Report says that there is a 1% to 10% chance that “CH₄ emissions from clathrates will substantially warm the climate system over the next few centuries”.³⁴⁵ It is unclear how to interpret this because it is not clear what they mean by ‘substantial’.

Sayedi et al (2020) carried out an expert elicitation study of 25 permafrost researchers on the stocks and sensitivity of carbon in the subsea permafrost domain. The study’s findings on cumulative emissions from the permafrost domain on different emissions scenarios are

³⁴¹ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 5.1.2.1.

³⁴² See section 3.3.2.

³⁴³ Frans-Jan W. Parmentier and Torben R. Christensen, “Arctic: Speed of Methane Release,” Comments and Opinion, *Nature*, August 28, 2013, <https://doi.org/10.1038/500529a>; Dirk Notz, Victor Brovkin, and Martin Heimann, “Arctic: Uncertainties in Methane Link,” Comments and Opinion, *Nature*, August 28, 2013, <https://doi.org/10.1038/500529b>.

³⁴⁴ Schuur et al., “Climate Change and the Permafrost Carbon Feedback.”

³⁴⁵ IPCC, *Climate Change 2021: The Physical Science Basis*, sec. 5.4.9.1.3.

shown below. The questions assume that concentrations reach a peak near to 2100 and then remain at that level indefinitely.³⁴⁶ As discussed in Chapter 2, in practice it is more likely that concentrations would slowly decline.

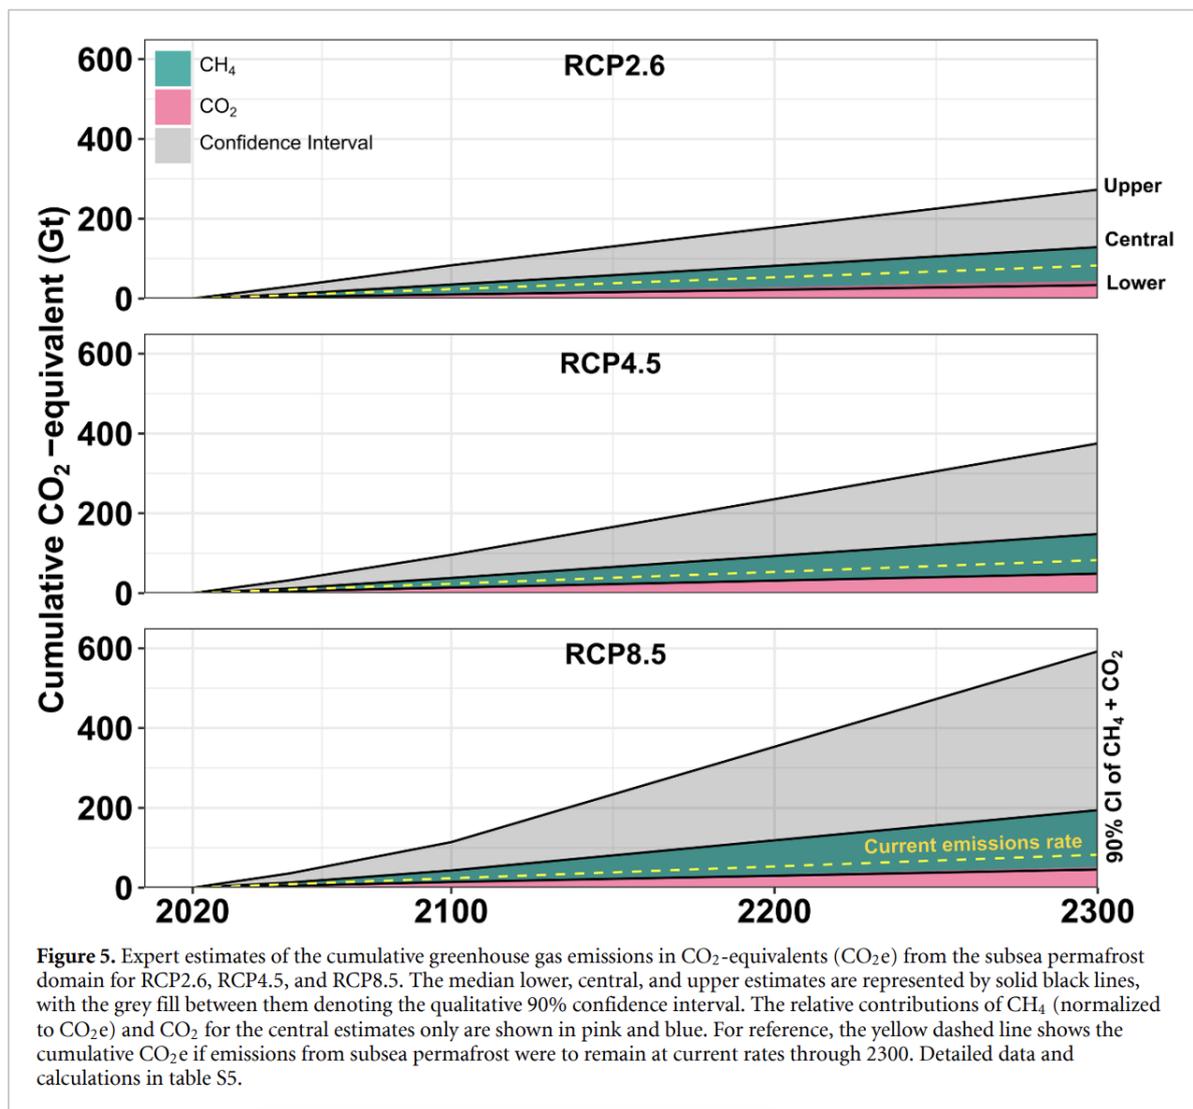


Figure 5. Expert estimates of the cumulative greenhouse gas emissions in CO₂-equivalents (CO₂e) from the subsea permafrost domain for RCP2.6, RCP4.5, and RCP8.5. The median lower, central, and upper estimates are represented by solid black lines, with the grey fill between them denoting the qualitative 90% confidence interval. The relative contributions of CH₄ (normalized to CO₂e) and CO₂ for the central estimates only are shown in pink and blue. For reference, the yellow dashed line shows the cumulative CO₂e if emissions from subsea permafrost were to remain at current rates through 2300. Detailed data and calculations in table S5.

Source: Sayedeh Sara Sayedi et al., ‘Subsea Permafrost Carbon Stocks and Climate Change Sensitivity Estimated by Expert Assessment’, *Environmental Research Letters* 15, no. 12 (December 2020): 5, <https://doi.org/10.1088/1748-9326/abcc29>.

The implications for warming of emissions from the permafrost region are summarised in the guesstimate figure below (from [this model](#)).

³⁴⁶ “The selected RCPs were RCP2.6, which has a peak concentration of ~490 ppm CO₂-equivalent (CO₂e) reached before 2100, RCP4.5 with a peak of ~650 ppm CO₂e at 2100, and RCP8.5 with a peak of ~1400 ppm CO₂e at 2100 (Moss et al 2010, Koenigk et al 2013).” Sayedeh Sara Sayedi et al., ‘Subsea Permafrost Carbon Stocks and Climate Change Sensitivity Estimated by Expert Assessment’, *Environmental Research Letters* 15, no. 12 (December 2020): 5, <https://doi.org/10.1088/1748-9326/abcc29>.

TCRE	1.5 0.83 to 2.5	Emissions 2100	Warming 2100	Emissions 2300	Warming 2300
RCP2.6		36 11 to 93	0.053 0.013 to 0.15	120 38 to 290	0.18 0.046 to 0.51
RCP4.5		44 16 to 100	0.067 0.019 to 0.17	170 52 to 430	0.25 0.065 to 0.7
RCP8.5		47 15 to 120	0.071 0.018 to 0.19	220 51 to 630	0.33 0.062 to 1.1

As mentioned in Chapter 2, the IPCC does not give a 90% confidence interval for the Transient Climate Response to Cumulative Emissions, but does give a 66% confidence interval of 1.0°C to 2.3°C per trillion tonnes of carbon. In the model, I have roughly guessed at a 90% confidence interval of 0.8 to 2.5°C per trillion tonnes of carbon

One concern with these numbers is that, because they rely on the CO₂-equivalent metric, they overweight the warming effect of methane.³⁴⁷ Thus, the estimates in the table above are likely biased high to some extent.

Nonetheless, even if the CO₂-equivalent metric is correct, this does illustrate that the claims of Whiteman et al (2013) about the likelihood of a sudden release of methane from clathrates causing 2°C of warming is well outside the scientific mainstream: on business as usual (RCP4.5), the 95th percentile of warming due to *both* permafrost carbon and methane clathrates is around 0.2°C by 2100.

8.3. Amazon forest dieback

Climate change has competing effects on forests. On the one hand, higher temperatures and increased drying is damaging to forests, but on the other hand higher CO₂ levels are good because of the CO₂ fertilisation effect, which encourages photosynthesis and the efficient use of water. Some early climate models found that global warming of 3°C would cause the Amazon rainforest to die off, which would in turn release further CO₂ into the atmosphere.³⁴⁸ The models produced this result even without taking account of deforestation and fires.

The researchers who initially found the effect suggest that the risk is probably smaller than first estimated.³⁴⁹

³⁴⁷ Myles Allen, 'Short-Lived Promise: The Science and Policy of Cumulative and Short-Lived Climate Pollutants', Oxford Martin School, 2015.

³⁴⁸ P. M. Cox et al., 'Amazonian Forest Dieback under Climate-Carbon Cycle Projections for the 21st Century', *Theoretical and Applied Climatology* 78, no. 1 (1 June 2004): 137–56, <https://doi.org/10.1007/s00704-004-0049-4>.

³⁴⁹ "The latest Earth system models show limited evidence of Amazon forest dieback in the absence of direct human deforestation. A handful of models show reductions in Amazonian forest cover due to climate change, but most models show increasing forest cover due to CO₂ fertilisation. None of these models include phosphorus limitations or forest fires, though, so the jury is still out. However, personally, I am less concerned about climate-change driven Amazon dieback than I was when we published our study in 2000." [Guest post](#) by Peter Cox on Carbon Brief.

Most Earth System Models now suggest that the effect of CO₂ fertilisation overwhelms the effect of climate change. However, there are reasons to think that the models could be biased in either direction.

- In the real world, plants might be less vulnerable than models suggest due to greater plant trait diversity and possible acclimation to warming.
- Plants might be more vulnerable than models suggest due to insect outbreaks, or model limitations in representing wildfires and droughts.³⁵⁰

The paleoclimate record provides some comfort on the risk of Amazon dieback. Temperatures were more than 10°C higher during the Eocene, but tropical forests flourished.³⁵¹ Willis and MacDonald posit that models project more damaging effects because they neglect CO₂ fertilisation.

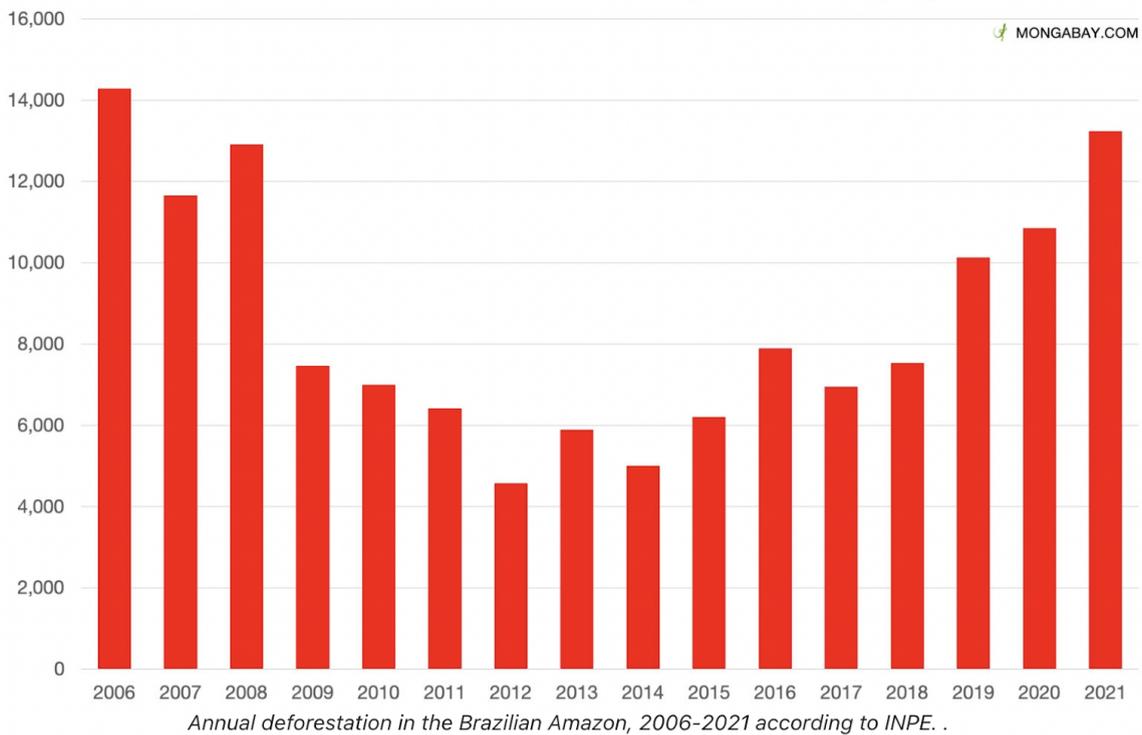
However, one important difference for future warming is that warming may occur in the context of rising deforestation.³⁵² Deforestation has been on the rise since 2012 and has accelerated since Bolsonaro's election in 2019. Amazonian deforestation in 2021 was the highest it has been since 2006.

³⁵⁰ "Most ESMs project continuing carbon accumulation in tropical forests as a result of direct CO₂ effects overwhelming the negative effects of climate change (Huntingford et al., 2013; Driifhout et al., 2015; Boulton et al., 2017). In the real world, forests may be less vulnerable to climate changes than those modelled in ESMs because of the greater plant trait diversity which confers additional resilience (Reyer et al., 2015; Levine et al., 2016; Sakschewski et al., 2016), and also because of possible acclimation of vegetation to warming (Good et al., 2011, 2013; Lloret et al., 2012; Mercado et al., 2018). Contrary, forests may be more vulnerable in the real world due to indirect climate change effects such as insect outbreaks and diseases not considered here (Section 5.4.3.2) or model limitations in representing the effects disturbances such as wildfire and droughts. In general, forests are most vulnerable when climate change is combined with increased rates of direct deforestation (Nobre et al., 2016; Le Page et al., 2017)." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Ch. 5, sec. 5.4.9.1.1.

³⁵¹ K. J. Willis and G. M. MacDonald, 'Long-Term Ecological Records and Their Relevance to Climate Change Predictions for a Warmer World', *Annual Review of Ecology, Evolution, and Systematics* 42, no. 1 (2011): sec. 2, <https://doi.org/10.1146/annurev-ecolsys-102209-144704>.

³⁵² Thanks to James Ozden for pressing me on this point.

Deforestation in the Brazilian Amazon, 2006-2015 (sq km)



Source: [Mongabay](https://www.mongabay.com)

Boulton et al (2022) point out that due to climate change and land use change, three quarters of the Amazon has been losing resilience since 2000.³⁵³ Thus, once we adjust for deforestation, Earth System Models likely understate the risks of dieback.

The IPCC Sixth Assessment Report estimates that on the pessimistic assumption (1-10% chance) of no CO₂ fertilisation, this would release 50GtC per 1°C of warming. So for 4°C, Amazon dieback would add at most 200 billion tonnes of carbon to the atmosphere,³⁵⁴ which would cause around 0.3°C of warming.

For more, see [this CarbonBrief post](#).

³⁵³ Chris A. Boulton, Timothy M. Lenton, and Niklas Boers, 'Pronounced Loss of Amazon Rainforest Resilience since the Early 2000s', *Nature Climate Change* 12, no. 3 (March 2022): 271–78, <https://doi.org/10.1038/s41558-022-01287-8>.

³⁵⁴ "In order to estimate an upper limit on the impact of Amazon forest dieback on atmospheric CO₂, we consider the very unlikely limiting case of negligible direct-CO₂ effects (Section 5.4.1). Emergent constraint approaches (Section 5.4.6) may be used to estimate an overall loss of tropical land carbon due to climate change alone, of around 50 PgC per °C of tropical warming (Cox et al., 2013b; Wenzel et al., 2014). This implies an upper limit to the release of tropical land carbon of <200 PgC over the 21st century (assuming tropical warming of <4°C, and no CO₂-fertilisation), which translates to dCO₂/dt < 0.5 ppm yr⁻¹" IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Ch. 5, sec. 5.4.9.1.1.

8.4. Boreal forest dieback

Boreal forests are found in the cold climates of the northern hemisphere high latitudes.³⁵⁵ They lie just to the south of the Arctic tundra, where tree growth is restricted by year-round freezing or near-freezing temperatures and a lack of rain. Boreal forests are characterised by species that can cope with the cold, such as pine, spruce and larch. They cover vast stretches of North America and northern Europe and Asia. They are the largest biome anywhere on Earth and account for 30% of the world's forests, and hold around a third of the world's terrestrial carbon.

The boreal region is warming twice as fast as the global average. Some models suggest that the parts of boreal forests in the south could pass a tipping point, and transition to open woodland and grassland, while the northern part of the forest would start to encroach into the Arctic permafrost. Due to the competing effects in the north and south, the IPCC claims that boreal forest is not expected to be a major source of CO₂ emissions.

“Boreal forest dieback is not expected to change the atmospheric CO₂ concentration substantially because forest loss at the south is partly compensated by (i) temperate forest invasion into the previous boreal area and (ii) boreal forest gain at the north (Friend et al., 2014; Kicklighter et al., 2014; Schaphoff et al., 2016) (medium confidence). An upper estimate of this magnitude, based on statistical modelling of climate change alone, is of 27 Pg vegetation C loss in the southern boreal forest, which is roughly balanced by gains in the northern zone (Koven, 2013). Carbon release from vegetation and soil due to wildfires in boreal regions (Eliseev et al., 2014b; Turetsky et al., 2015; Walker et al., 2019a) is also not expected to change this estimate substantially because of its small present-day value of about 0.2 PgC yr⁻¹ (van der Werf et al., 2017), and because of likely increases in precipitation in boreal regions.”³⁵⁶

8.5. Collapse of the Atlantic Meridional Overturning Circulation

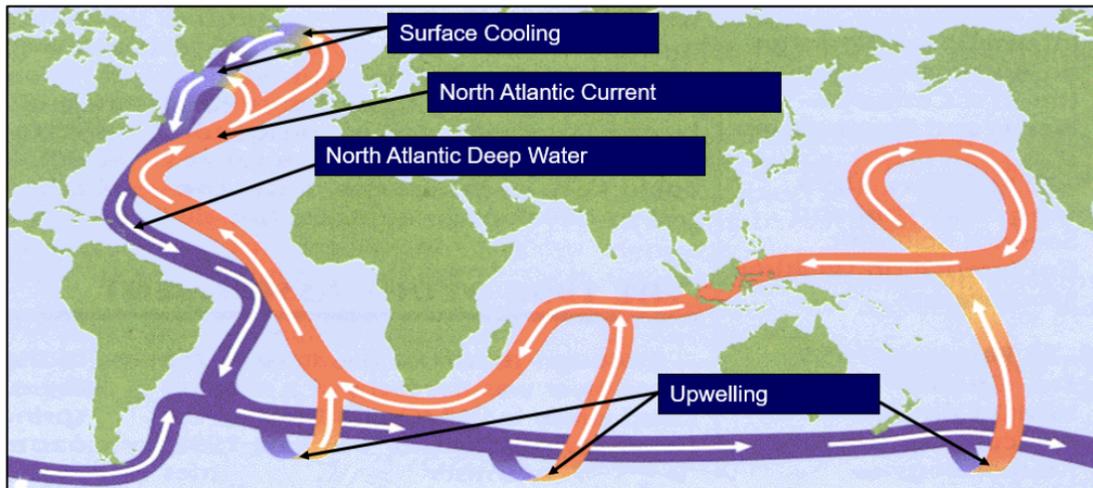
The Atlantic Meridional Overturning Circulation, known as the AMOC, is an ocean system that plays a major role in regulating the climate. It is driven by a delicate balance of ocean temperatures and salinity, which is at risk from being upset by a warming climate.

The potential collapse of the Atlantic Meridional Overturning Circulation is one of the more worrying tipping points discussed in the literature. Carbonbrief's review of AMOC collapse is available [here](#).

The diagram below shows a schematic of the AMOC

³⁵⁵ All of this is from [CarbonBrief](#).

³⁵⁶ IPCC, Climate Change 2021: The Physical Science Basis, sec. 5.4.9.1.1.



Schematic of the AMOC. The red pathways show warmer water nearer the surface, while the purple pathways show colder, more dense water moving at depth. Credit: Met Office

Climate change affects the AMOC by diluting the water at higher latitudes (through ice melt and rainfall) and heating it up. This reduces the density of the water and weakens the AMOC. Weakening of the AMOC could happen gradually or suddenly. The AMOC would only recover several centuries after emissions stop.

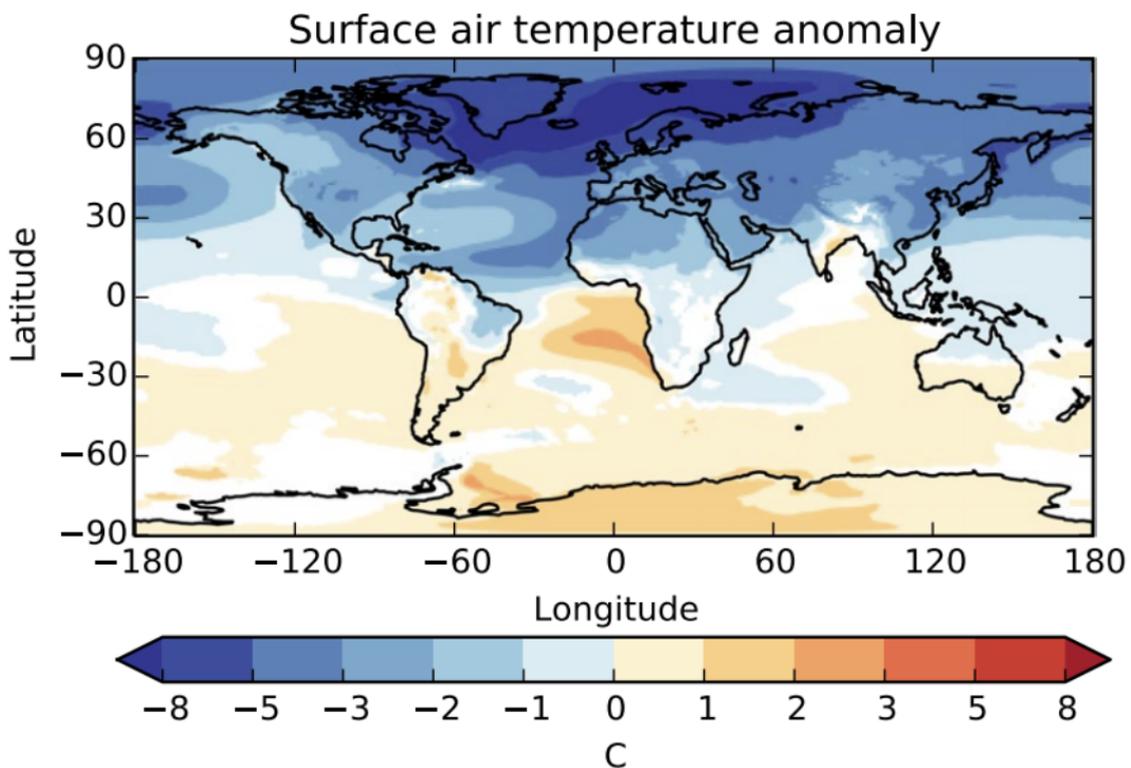
About 14,500 years ago, the Earth started to emerge from the Last Ice Age to the warmer Holocene interglacial. Partway through this transition, temperatures in the Northern Hemisphere suddenly returned to glacial conditions over the course of decades. This is known as the [Younger Dryas](#) event and it is thought to have been caused by a severe weakening of the AMOC due to an influx of fresh meltwater into the North Atlantic. The Younger Dryas also caused widespread changes in precipitation: the African and Asian monsoons weakened but those in the Southern Hemisphere strengthened.³⁵⁷

³⁵⁷ “During the last deglacial transition, one such slowdown in AMOC—during the Younger Dryas event (12,800–11,700 years ago)—caused worldwide changes in precipitation patterns. These included a southward migration of the tropical ITCZ (Peterson et al., 2000; McGee et al., 2014; Schneider et al., 2014; Mohtadi et al., 2016; Reimi and Marcantonio, 2016; Wang et al., 2017c) and systematic weakening of the African and Asian monsoons (Tierney and deMenocal, 2013; Otto-Bliesner et al., 2014; Cheng et al., 2016; Grandey et al., 2016; Wurtzel et al., 2018). Conversely, the Southern Hemisphere monsoon systems intensified (Cruz et al., 2005; Ayliffe et al., 2013; Strikis et al., 2015, 2018a; Campos et al., 2019). Drying occurred in Mesoamerica (Lachniet et al., 2013) while the North American monsoon system was largely unaffected (Bhattacharya et al., 2018). The mid-latitude region in North America was wetter (Polyak et al., 2004; Grimm et al., 2006; Wagner et al., 2010; Voelker et al., 2015), while Europe was drier (Genty et al., 2006; Rach et al., 2017; Naughton et al., 2019).” IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, sec. 8.6.1.

8.5.1. Climatic effects of AMOC collapse

A collapse of the AMOC would have a range of climatic effects on sea levels and floods,³⁵⁸ on winter storms,³⁵⁹ and on temperature and precipitation. I will focus on temperature and precipitation here.

The figure below illustrates the changes that result compared to the pre-collapsed state according to one model. Shutdown of the AMOC results in a cooling (blue shading) of the whole northern hemisphere, particularly the regions closest to the zone of North Atlantic heat loss (the “radiator” of the North Atlantic central heating system). In these regions the cooling exceeds the projected warming due to greenhouse gases, so a complete shutdown in the 21st century could result in a net cooling in regions such as western Europe.

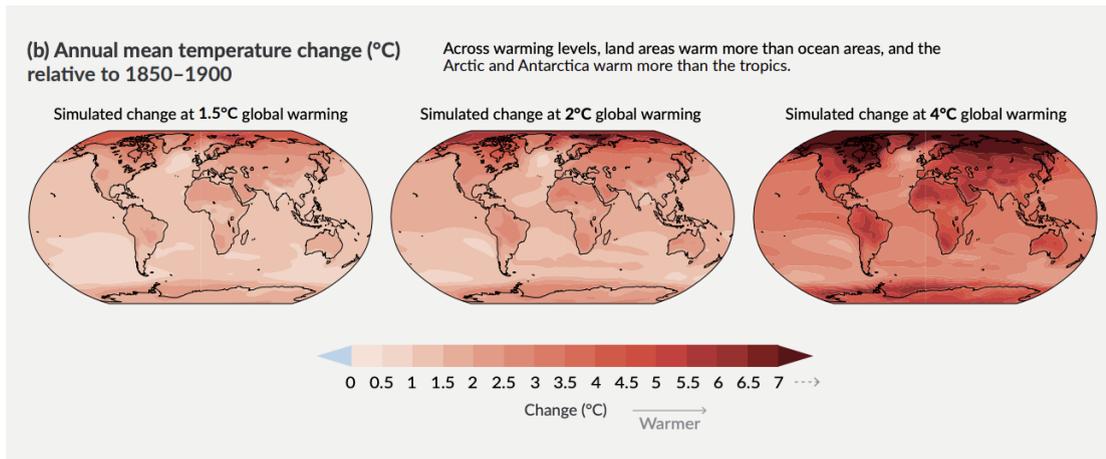


Modelled change in surface temperature (C) following an artificially induced collapse of the AMOC. Shading indicates cooling (blue) or warming (orange and red). Reprinted by permission from Springer. Jackson et al. (2015) Global and European climate impacts of a slowdown of the AMOC in a high resolution GCM, Climate Dynamics.

For reference, the map below from the IPCC shows future regional warming at different levels of global warming:

³⁵⁸ Christopher M. Little et al., 'The Relationship Between U.S. East Coast Sea Level and the Atlantic Meridional Overturning Circulation: A Review', *Journal of Geophysical Research: Oceans* 124, no. 9 (2019): 6435–58, <https://doi.org/10.1029/2019JC015152>.

³⁵⁹ T. Woollings et al., 'Response of the North Atlantic Storm Track to Climate Change Shaped by Ocean–Atmosphere Coupling', *Nature Geoscience* 5, no. 5 (May 2012): 313–17, <https://doi.org/10.1038/ngeo1438>.



Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Summary for Policymakers, SPM. 5.

This suggests that areas around the North Atlantic might cool by 1-2°C on net.

The IPCC has low confidence in model projections of the timing and magnitude of AMOC weakening in the 21st Century, and there is large disagreement across models about the regional effects.³⁶⁰

There would also be large effects on precipitation patterns. The pane on the left shows the effect on precipitation during the Younger Dryas, and the right pane shows the effect on precipitation due to modelled AMOC collapse.

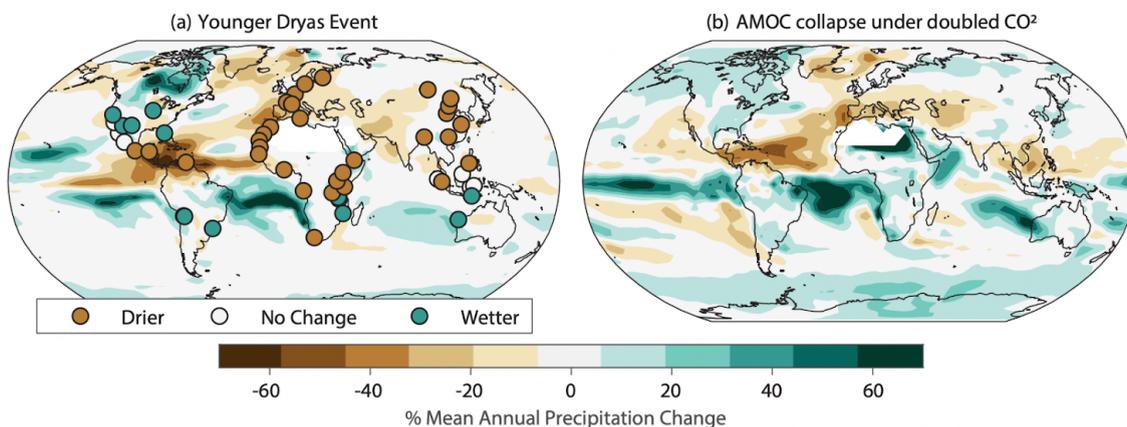


Figure 8.27 | Comparison of reconstructed past and idealized future annual mean precipitation responses to an Atlantic Meridional Overturning Circulation (AMOC) collapse. (a) Model simulation of precipitation response to the Younger Dryas event relative to the preceding warm Bölling-Allerød period (base colours, calculated as the difference between 12,600–11,700 years before the present (BP) and 14,500–12,900 BP from the Transient Climate Evolution (TraCE) paleoclimate simulation of Liu et al., 2009), with paleoclimate proxy evidence superimposed on top (dots). (b) Model simulation of precipitation response to an abrupt collapse in AMOC under a doubling of 1990 CO₂ levels (after W. Liu et al., 2017). Regions with rainfall rates below 20 mm yr⁻¹ are masked. Further details on data sources and processing are available in the chapter data table (Table 8.SM.1).

³⁶⁰ “Note that these ranges are based on ensemble means of individual models, largely smoothing out internal variability. If single realizations are considered, the ranges become wider, especially by lowering the low end of the range (Section 4.3.2.3). In summary, it is very likely that AMOC will decline in the 21st century, but there is low confidence in the model’s projected timing and magnitude.” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 9.2.3.1.

Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Chap 8, Fig. 8.27.

Since there is so much uncertainty across models, the true net effect may be significantly different.

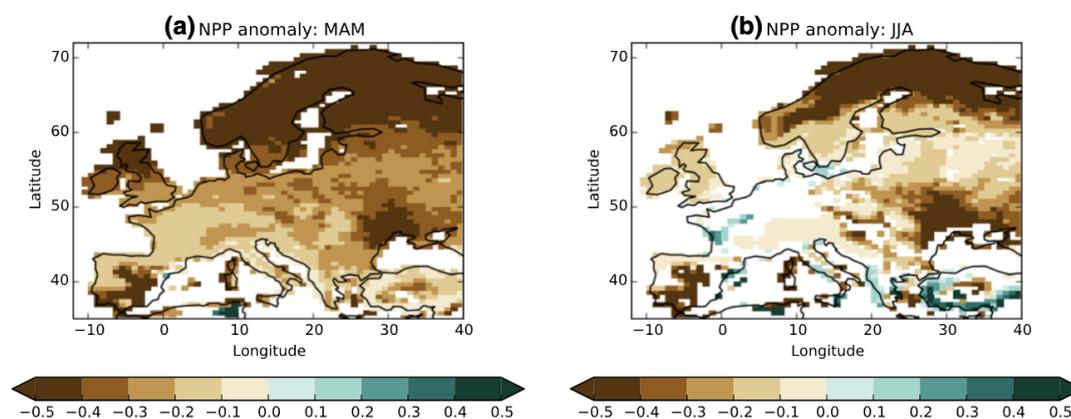
AMOC collapse would also affect global monsoons, with African and Asian monsoons weakening and southern hemisphere monsoons strengthening.³⁶¹

8.5.2. How would AMOC collapse affect human society?

The collapse of the AMOC would have a range of negative effects on society. In Northern Europe, winters would be colder, which would likely increase temperature-related deaths.³⁶² Winter storminess could also cause damage to infrastructure. Probably the most damaging effects would occur due to changes to monsoons in Africa and Asia. This could have very bad effects by creating droughts and damaging agriculture.

Cooling is generally worse for agriculture than warming because one day of frost destroys the entire growing season: this is why nuclear winter is so bad. Reduced precipitation and soil moisture is bad for obvious reasons.

As far as I am aware, there have been no studies on the effects that AMOC collapse might have on global agriculture. Jackson et al (2015) explore the effect on net primary productivity of vegetation in Western Europe in spring (left pane) and summer (right pane):



Source: Jackson et al., 'Global and European Climate Impacts of a Slowdown of the AMOC in a High Resolution GCM', Fig. 13.

Some regions suffer net primary productivity losses of 50%, though the effects are more modest in many regions.

Ritchie et al (2020) explore the effects of AMOC collapse on farming in Britain. They conclude that in the absence of irrigation:

³⁶¹ "A collapse of the Atlantic Meridional Overturning Circulation could weaken the African and Asian monsoons but strengthen the Southern Hemisphere monsoons (high confidence)." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box TS. 13.

³⁶² See Chapter 6.

“The expected overall area of arable production is predicted to fall dramatically from 32 to 7% of land area (Extended Data Figs. 2 and 3). This in turn generates a major reduction in the value of agricultural output, with a decrease of £346 million per annum (Table 1), representing a reduction in total income from British farming of ~10% (ref. 33). The key driver of the arable loss seen across Great Britain is climate drying due to AMOC collapse, rather than cooling.”³⁶³

With irrigation:

“land area under arable production still increases from 32 to 38% by 2080, with an accompanying increase in output value of £79 million per annum”³⁶⁴

The irrigation would cost more than £800 million per year, and would be technically difficult. UK GDP is currently around £2.2 trillion, so this would be 0.3% of UK GDP. We should expect crop yields to have improved substantially in the next few decades due to technological improvements. For reference, over the last 60 years, crop yields have increased by upwards of 200%. For the purposes of producing enough food to feed its own people, the UK would be rich enough to adapt, though the economic costs would be large.

The cooling effects of AMOC collapse would not be sufficient to shorten the growing season in other regions because any cooling effects would be cancelled out by higher background temperatures due to the strengthening greenhouse effect.

The effects of precipitation changes are also harder to predict. Irrigation would be an effective response to any shortfall of rain, but there would be large transition costs for affected countries. For countries that are poor and agrarian at the time of AMOC collapse, the scope for adaptation would be much more limited. AMOC collapse is expected to weaken the monsoons around West Africa, the Sahel and India, which supply much of the annual rainfall in those regions. As I discussed above, weakening of the West African monsoon between the late 1960s and 1980s caused a drought that killed tens of people.

Overall, it is clear that global agriculture would still be viable, and that rich countries that are rich at the time of collapse would be able to adapt to AMOC collapse: people would not starve, though they would suffer substantial economic costs. Some regions would lose out while others would gain, though the net effect seems likely to be strongly negative given the populations in losing regions (India, West Africa and the Sahel).

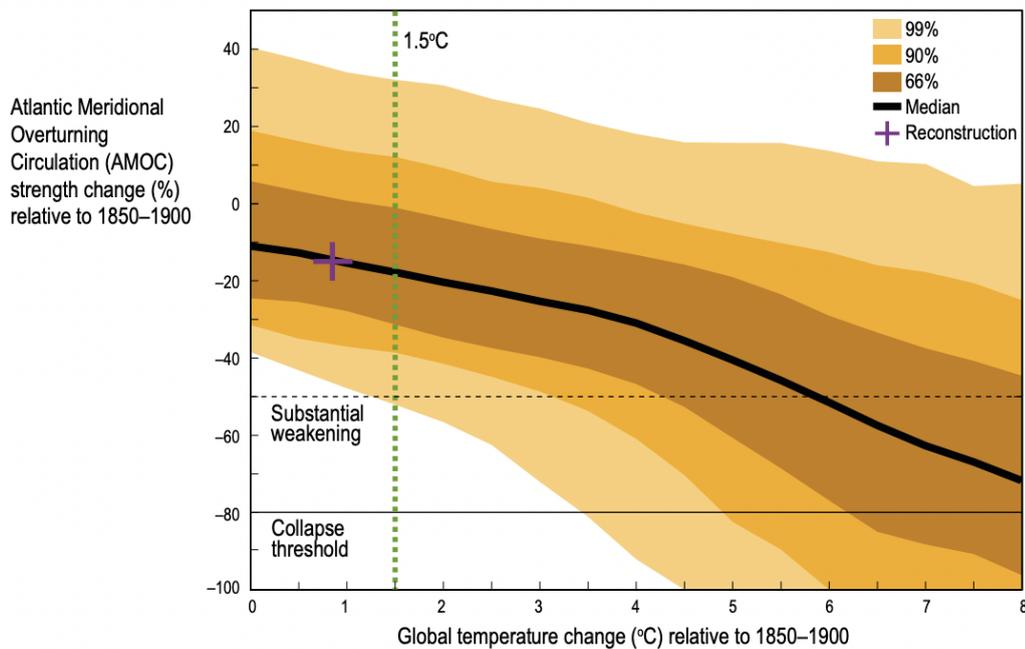
8.5.3. The probability of AMOC collapse

The AMOC is projected to weaken on all emissions scenarios up to 2100. The IPCC's Special Report on the Ocean and the Cryosphere in a Changing Climate claimed that AMOC

³⁶³ Paul DL Ritchie et al., 'Shifts in National Land Use and Food Production in Great Britain after a Climate Tipping Point', *Nature Food* 1, no. 1 (2020): 77.

³⁶⁴ Paul DL Ritchie et al., 'Shifts in National Land Use and Food Production in Great Britain after a Climate Tipping Point', *Nature Food* 1, no. 1 (2020): 79.

collapse was about as likely as not by 2300.³⁶⁵ The IPCC provides the following chart which shows the probability of AMOC collapse at different levels of warming.



Source: IPCC, 'Special Report on the Ocean and Cryosphere in a Changing Climate', 2019, Fig 6.9.

As this shows, at 4°C of warming, the risk of AMOC collapse is around 5%, rising to nearly 50% at 8°C. However, the IPCC cautions that existing models do not capture some processes that are relevant to the risk of AMOC collapse, so the above chart probably understates the risk.³⁶⁶

The AMOC would only recover several centuries after the cessation of emissions.³⁶⁷

³⁶⁵ "Both AR5 (Collins et al., 2013) and SROCC (Collins et al., 2019) assessed that an abrupt collapse of AMOC before 2100 was very unlikely, but SROCC added that, by 2300, an AMOC collapse was as likely as not for high-emissions scenarios." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 9.2.3.1.

³⁶⁶ "The SROCC also assessed that model-bias may considerably affect the sensitivity of the modelled AMOC to freshwater forcing. Tuning towards stability and model biases (Valdes, 2011; Liu et al., 2017; Mecking et al., 2017; Weijer et al., 2019) provides CMIP models a tendency toward unrealistic stability (medium confidence). By correcting for existing salinity biases, Liu et al. (2017) demonstrated that AMOC behaviour may change dramatically on centennial to millennial timescales and that the probability of a collapsed state increases. None of the CMIP6 models features an abrupt AMOC collapse in the 21st century, but they neglect meltwater release from the Greenland ice sheet and a recent process study reveals that a collapse of the AMOC can be induced even by small-amplitude changes in freshwater forcing (Lohmann and Ditlevsen, 2021). As a result, we change the assessment of an abrupt collapse before 2100 to medium confidence that it will not occur" IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 9.2.3.1.

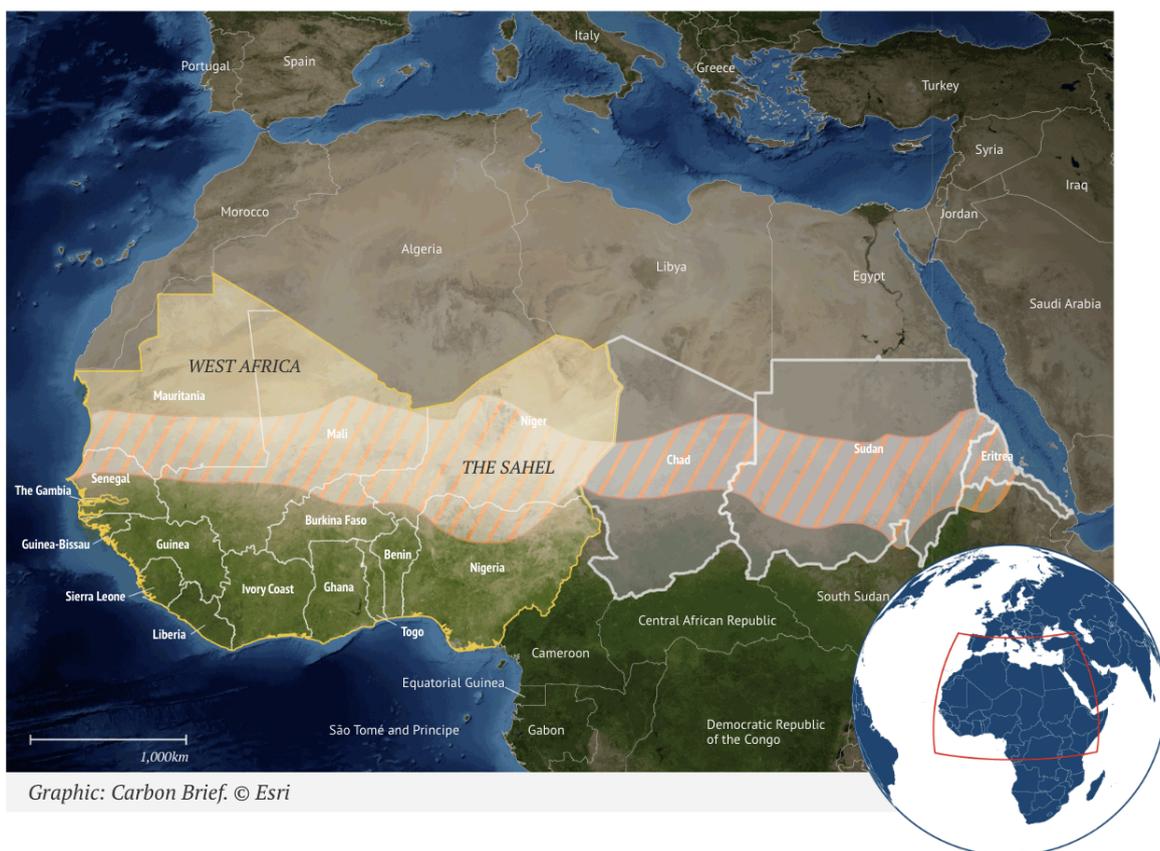
³⁶⁷ "Based on a large initial condition ensemble of simulations with a CMIP5 model (CanESM2) with emission scenarios leading to stabilization of global warming of 1.5°C, 2.0°C, or 3.0°C relative to 1850–1900, AMOC continues to decline for 5–10 years after GSAT is effectively stabilized at the given GWL (Sigmond et al., 2020). This is followed by a recovery of AMOC strength for about the next 150 years to a level that is approximately independent of the considered stabilization scenario. These results are replicated in simulations in a CMIP6 model (CanESM5) with emissions cessation after diagnosed CO2 emissions reach 750 Gt, 1000 Gt, or 1500 Gt. These emissions levels lead to global

8.6. Changes to the West African monsoon

The Carbonbrief review of the West African monsoon is [here](#). My introduction here borrows from the Carbonbrief review.

The term “monsoon” in its strictest sense refers to the seasonal reversal of winds and its accompanying rainfall. Along with India, West Africa is one of the few places on Earth where this happens.

The West African monsoon brings rainfall to West Africa and the Sahel – a band of semi-arid grassland sandwiched between the Sahara desert to the north and tropical rainforests to the south. The Sahel stretches from the Atlantic coast of Mauritania and Senegal through to Sudan, Eritrea and the Red Sea.



Graphic: Carbon Brief. © Esri

West Africa’s dry season, which runs from November through to May, sees prevailing dry and dusty winds come from the desert. The monsoon brings rain to the region from around June to September.

The West African monsoon is notoriously unreliable. Between the late 1960s and 1980s, a lack of rain hit much of the Sahel, with average rainfall declining by more than 30% over most of the region compared to the 1950s. This plunged the region into an extended

warming stabilization at 1.5°C, 2.0°C, or 3.0°C relative to 1850–1900. In summary, in these model simulations the AMOC recovers over several centuries after the cessation of CO₂ emissions.” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 4.3.2.3.

drought, contributing to a famine that killed tens of thousands of people and triggering an international aid effort.

Models and paleoclimate evidence suggests that greenhouse warming will strengthen monsoon precipitation. However, anthropogenic aerosol emissions in the second half of the 20th Century outweighed the effect of greenhouse gases, causing drying.³⁶⁸

Paleoclimate evidence suggests that there were past 'Green Sahel' states in which rainfall was much higher in the Sahel.³⁶⁹ Models suggest that climate change will have mixed effects on rainfall patterns in the Sahel, with monsoon precipitation increasing in the central Sahel, but decreasing in the Western Sahel.³⁷⁰

All of the evidence suggests that the Sahel is very sensitive to climatic drivers and that there could be abrupt shifts in precipitation.

As discussed above, the West African monsoon would also be affected by AMOC collapse.

Impacts on the monsoon are reversible within years to decades if greenhouse gas concentrations are reduced.³⁷¹

8.7. Indian monsoon shift

The Carbonbrief review of Indian monsoon shift is [here](#). The introduction here borrows from the Carbonbrief review.

India receives around 70% of its annual rainfall during the monsoon season. For some areas of western and central India, it accounts for as much as 90%. The monsoon rains are crucial for India's farm sector, which makes up about a sixth of India's economy and employs about half of the country's 1.3 billion population. The monsoon season starts around June and typically ends at the end of August.

³⁶⁸ "Northern Hemispheric anthropogenic aerosols weakened the regional monsoon circulations in South Asia, East Asia and West Africa during the second half of the 20th century, thereby offsetting the expected strengthening of monsoon precipitation in response to GHG-induced warming (high confidence)." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box TS.13.

³⁶⁹ "Paleoclimate reconstructions provide evidence of past Green Sahara states (DeMenocal and Tierney, 2012), under which rainfall rates increased by an order of magnitude (Tierney et al., 2017), leading to a vegetated landscape (Jolly et al., 1998) with large lake basins (Gasse, 2000; Drake and Bristow, 2006)." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Ch. 8, sec. 8.6.2.2.

³⁷⁰ "Over South and South East Asia, East Asia and the central Sahel, monsoon precipitation is projected to increase, whereas over North America and the far western Sahel it is projected to decrease (medium confidence)." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box TS.13.

³⁷¹ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Chap 4, Table 4.10.

The IPCC has medium confidence that climate change will increase the strength of the Indian monsoon.³⁷² Monsoons are inherently hard to model, and there is scientific disagreement about how climate change will affect the Indian monsoon.³⁷³

Some studies suggest that there could be an abrupt transition from wet to dry monsoon over the course of hundreds to thousands of years if there is high air pollution or low greenhouse gas levels or both.³⁷⁴

As discussed above, the Indian monsoon would also be affected by AMOC collapse.

Impacts on the monsoon are reversible within years to decades if greenhouse gas concentrations are reduced.³⁷⁵

8.8. Cloud feedbacks

In my view, the most worrying potential non-linearities in the climate system stem from cloud feedbacks. Most of the uncertainty about climate sensitivity is driven by uncertainty about clouds.³⁷⁶ There are several reasons to think that clouds may contribute to non-linear increases in warming as emissions and warming increase.

Firstly, there is evidence from the paleoclimate that cloud feedbacks may account for the fact that equilibrium climate sensitivity seems to rise with temperatures, or is 'state-dependent'.³⁷⁷ The early Eocene period, 56 to 48 million years ago, was the hottest period in the Cenozoic. During the Paleocene-Eocene Thermal Maximum, temperatures increased by 5°C following an increase of atmospheric CO₂ of 70% to 100%, which suggests an equilibrium climate

³⁷² "For the North American monsoon, projections indicate a decrease in precipitation, whereas increased monsoon rainfall is projected over South and South East Asia and over East Asia (medium confidence) (Box TS.13, Figure 1)." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Box TS.13.

³⁷³ See the discussion in the Carbonbrief [overview](#).

³⁷⁴ "Some papers have suggested the possibility for more abrupt changes in the Indian monsoon. A 2005 study, for example, used a simple model to identify the possibility of the monsoon having two stable states: wet (as it is now) and second state characterised by low rainfall. Key to these two states is the so-called "moisture advection feedback". This, the paper explains, is where "the land to ocean pressure gradient, which drives the monsoon circulation, is reinforced by the moisture the monsoon itself carries from the adjacent Indian Ocean". In other words, a significant factor in maintaining the monsoon is the heat released when the water vapour it holds condenses to form rain. Another paper, published in the Proceedings of the National Academy of Sciences (PNAS) in 2009, suggests this feedback acts as an "internal amplifier" for the monsoon. The implication is that this feedback magnifies anything that affects the air pressure gradient generated by warm air rising over the Asian landmass. Thus "relatively weak external perturbations" could lead to "abrupt changes" in the monsoon, the PNAS paper says. The model simulations in the 2005 study suggest how a switch between states could be triggered. This includes cooling of the land surface through large amounts of air pollution, cooling through very low CO₂ levels in the atmosphere, or a combination of the two." See the discussion in the Carbonbrief [overview](#).

³⁷⁵ Source: IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), Chap 4, Table 4.10.

³⁷⁶ Bjorn Stevens and Sandrine Bony, 'What Are Climate Models Missing?', *Science* 340, no. 6136 (31 May 2013): 1053–54, <https://doi.org/10.1126/science.1237554>.

³⁷⁷ Jonah Bloch-Johnson et al., 'Climate Sensitivity Increases Under Higher CO₂ Levels Due to Feedback Temperature Dependence', *Geophysical Research Letters* 48, no. 4 (2021): e2020GL089074, <https://doi.org/10.1029/2020GL089074>.

sensitivity of 5°C, well in excess of IPCC estimates.³⁷⁸ For this reason, climate models without state-dependent climate sensitivity have struggled to reproduce the Eocene climate.³⁷⁹

This is how climate sensitivity changed throughout the Eocene period from 56 million years ago to 33 million years ago.

³⁷⁸ Jiang Zhu, Christopher J. Poulsen, and Jessica E. Tierney, 'Simulation of Eocene Extreme Warmth and High Climate Sensitivity through Cloud Feedbacks', *Science Advances* 5, no. 9 (1 September 2019): eaax1874, <https://doi.org/10.1126/sciadv.aax1874>.

³⁷⁹ "The Early Eocene, a period of elevated atmospheric CO₂ (>1000 ppmv), is considered an analog for future climate. Previous modeling attempts have been unable to reproduce major features of Eocene climate indicated by proxy data without substantial modification to the model physics." Jiang Zhu, Christopher J. Poulsen, and Jessica E. Tierney, 'Simulation of Eocene Extreme Warmth and High Climate Sensitivity through Cloud Feedbacks', *Science Advances* 5, no. 9 (1 September 2019): eaax1874, <https://doi.org/10.1126/sciadv.aax1874>.

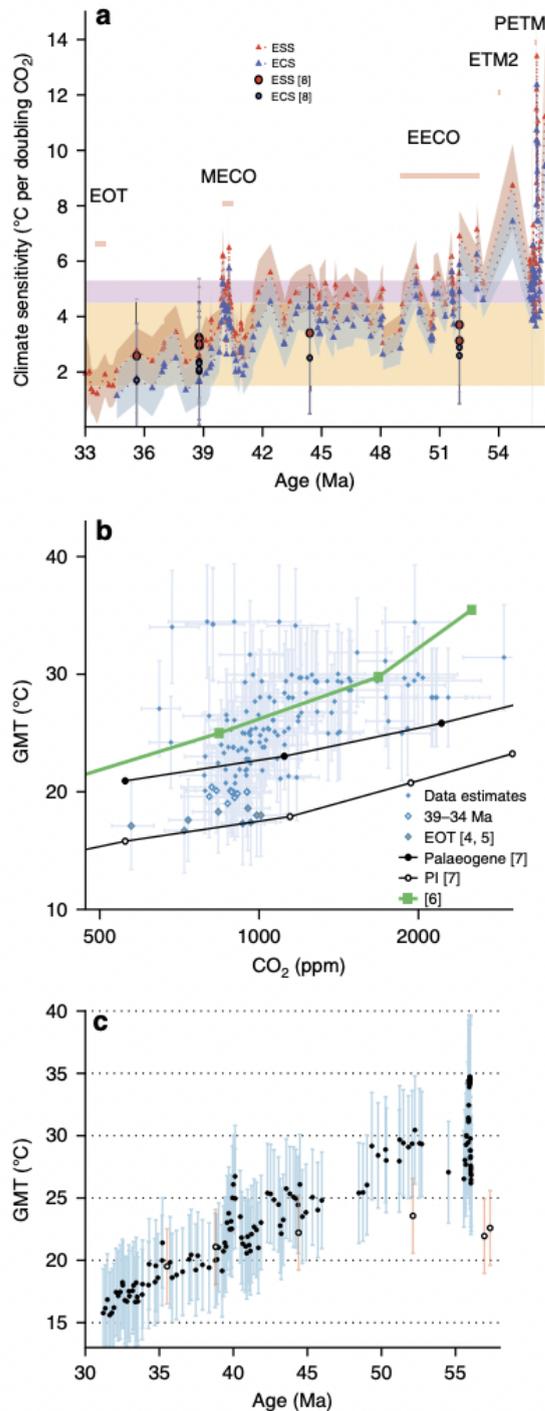


Fig. 5 Evolving climate sensitivity for the Eocene. **a** Calculated ESS (red triangles and error envelope), and ECS (blue triangles and error envelope). See text for relevant methodology. Orange area represents the IPCC range in ECS⁹⁵, and the pink highlighted area the updated 20th century ECS with the addition of state-of-the-art cloud physics⁹⁶. Circles represent estimates from ref. ⁸. **b** Data-model inter-comparison, with all diamonds representing data. Open diamonds are the data between 39 and 34 Ma, and orange filled diamonds the EOT. Circles⁷ and squares⁶ are all model derived relationships (PI preindustrial). Uncertainties and error envelopes represent 1 s.d. of Monte Carlo propagated uncertainties. **c** Evolving GMT relationship for the Eocene. GMT is calculated using the BAYSPAR TEX₈₆ record from ODP 959. Error bars represent the calibration and analytical uncertainty on TEX₈₆. For comparison, the GMT estimates from ref. ⁸ are presented with open symbols and red error bars. NB the elevated ECS early in the PETM and ETM 2 are most likely a consequence of slight age model misalignments, or imply non-CO₂ forcing early in these events.

Source: E. Anagnostou et al., 'Proxy Evidence for State-Dependence of Climate Sensitivity in the Eocene Greenhouse', *Nature Communications* 11, no. 1 (7 September 2020): 4436, <https://doi.org/10.1038/s41467-020-17887-x>.

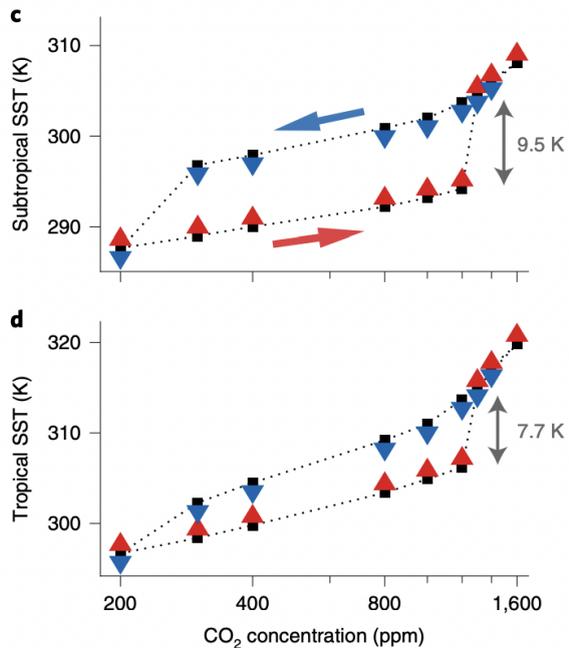
Newer climate models that better represent possible nonlinear cloud feedbacks suggest that equilibrium climate sensitivity is state-dependent. These models can in turn better represent the climate of past warm periods like the Eocene.³⁸⁰

Goldblatt et al (2021) argue that cloud feedbacks may explain the 'Faint Young Sun Problem', which is that the Sun was dimmer earlier in Earth's history, but glaciation was rare in the Precambrian. According to their model, because the Sun was dimmer and greenhouse gas levels were high in the precambrian, there was a substantial decrease in stratocumulus decks and consequently a decrease in planetary reflectivity.³⁸¹

While it is generally agreed that cloud feedbacks could have non-linear effects on temperature, there is disagreement about the nature and magnitude of the effect. Some models suggest that the warming from cloud feedbacks will be fairly smooth and gradual, whereas others project effects that are huge and sudden. Perhaps the most worrying climate change paper in recent years is Schneider et al (2019), which suggest that once CO₂ concentrations pass 1,200-1,300ppm, there would be 8°C of global warming over the course of days, on top of the 5°C warming we would already have lived through. In total, there would be around 13°C of warming relative to pre-industrial levels.

³⁸⁰ "Overall, the modelling evidence indicates that there is medium confidence that the net feedback parameter, α , increases (i.e., becomes less negative) with increasing temperature (i.e., that sensitivity to forcing increases with increasing temperature), under global surface background temperatures at least up to 40° C (Meraner et al., 2013; Seeley and Jeevanjee, 2021), and medium confidence that this temperature dependence primarily derives from increases in the water vapour and shortwave cloud feedbacks. This assessment is further supported by recent analysis of CMIP6 model simulations (Bloch-Johnson et al., 2020) in the framework of nonlinMIP (Good et al., 2016), which showed that out of ten CMIP6 models, seven of them showed an increase of the net feedback parameter with temperature, primarily due to the water vapour feedback." IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 7.4.3.1.

³⁸¹ "We perform global climate model experiments, using two versions of the Community Atmosphere Model, in which a reduced solar constant is offset by higher CO₂. Model runs corresponding to past climate show a substantial decrease in low clouds and hence planetary albedo compared with present, which contributes 40% of the required forcing to offset the faint Sun. Through time, the climatically important stratocumulus decks have grown in response to a brightening Sun and decreasing greenhouse effect, driven by stronger cloud-top radiative cooling (which drives low cloud formation) and a stronger inversion (which sustains clouds against dry air entrainment from above). We find that systematic changes to low clouds have had a major role in stabilizing climate through Earth's history, which demonstrates the importance of physical feedbacks on long-term climate stabilization, and a smaller role for geochemical feedbacks." Colin Goldblatt, Victoria L. McDonald, and Kelly E. McCusker, 'Earth's Long-Term Climate Stabilized by Clouds', *Nature Geoscience* 14, no. 3 (March 2021): 143–50, <https://doi.org/10.1038/s41561-021-00691-7>.



Source: Tapio Schneider, Colleen M. Kaul, and Kyle G. Pressel, 'Possible Climate Transitions from Breakup of Stratocumulus Decks under Greenhouse Warming', *Nature Geoscience* 12, no. 3 (March 2019): 163–67, <https://doi.org/10.1038/s41561-019-0310-1>, Fig. 3.

Although the model finds that warming would happen over the course of around a month, in reality the transition is likely to take years to decades.³⁸² There is also uncertainty about when the threshold would be crossed, and alternative assumptions imply that the threshold could be between 1,400ppm to 2,200ppm.³⁸³

On RCP8.5, concentrations would reach around 900ppm in 2100. As I mentioned in Chapter 1, RCP8.5 is a very extreme scenario. On current policy (RCP4.5), concentrations would reach around 550ppm. Nonetheless, the threshold is within reach in principle: if we burn all of the recoverable fossil fuels - 3 trillion tonnes of carbon - concentrations would rise to around 1,600ppm.³⁸⁴

³⁸² "The breakup of the stratocumulus clouds is more rapid than it would be in nature because of the unrealistically small thermal inertia of the underlying slab ocean." Schneider, Kaul, and Pressel, 'Possible Climate Transitions from Breakup of Stratocumulus Decks under Greenhouse Warming', SI p. 5; Tapio Schneider, personal communication, 20th August 2021.

³⁸³ "The CO₂ level at which the instability occurs depends on how largescale dynamics change with climate, which is heuristically parameterized in our simulations and hence is uncertain. In particular, the large-scale subsidence in the troposphere weakens under warming³², which lifts the cloud tops and counteracts the instability^{15,19,24}. Indeed, when we weaken the parameterized large-scale subsidence by 1 or 3% per Kelvin of tropical SST increase (within the range of GCM responses to warming³³), the stratocumulus instability occurs at higher CO₂ levels: around 1,400ppm with 1%K⁻¹ subsidence weakening, and around 2,200ppm with 3%K⁻¹ (Fig. 4)." Schneider, Kaul, and Pressel, 'Possible Climate Transitions from Breakup of Stratocumulus Decks under Greenhouse Warming'

³⁸⁴ N. S. Lord et al., 'An Impulse Response Function for the "Long Tail" of Excess Atmospheric CO₂ in an Earth System Model', *Global Biogeochemical Cycles* 30, no. 1 (2016): fig. 2, <https://doi.org/10.1002/2014GB005074>.

8.8.1. How plausible are rapid cloud feedbacks?

This research is controversial, and scientists are divided on how plausible it is. In the discussion by [CarbonBrief](#), several scientists say that they think the result is plausible, whereas a news article in [Science](#) interviewed several scientists who were more sceptical.

Unfortunately, it is difficult to know just what would happen at 1,300ppm because CO₂ concentrations have not been that high for at least tens of millions of years. For the early Eocene, the best temperature proxies only have ~4,000 year resolution.³⁸⁵ Thus, the speed of a ~100-year feedback would not show up in the proxy record. The IPCC notes that there is little evidence of such extreme nonlinear warming in the paleoclimate record, possibly short of the Paleocene-Eocene Thermal Maximum, which started from a baseline 12°C warmer than pre-industrial.³⁸⁶ Overall, cloud feedback in some form seems like a plausible cause of state-dependent climate sensitivity that we see in the paleoclimate evidence, but it is hard to know just how rapid and drastic the feedback might be. We cannot rule out the feedback found in Schneider et al (2019).

Finally, it is interesting to explore the potential impact such a tipping point might have on ecosystems. If the tipping point were real, then we would expect it to have kicked in during the Paleocene-Eocene Thermal Maximum when temperatures were 17°C higher and CO₂ concentrations were more than 2,000ppm. But if there were 8°C of warming over the course of years, then according to many climate-biodiversity models, we should expect there to have been a huge extinction event. The warming would be so rapid that it would outpace the ability of species to shift their range to their new ecological niche. So, either: the tipping point is not real, or the biodiversity models are wrong. I argued in Chapter 3 that the biodiversity models are probably wrong: ecosystems seem quite resilient even to rapid warming. Schneider et al (2019) should update us further towards that position.

Given progress on climate policy, it now seems unlikely that we could reach the 1,200ppm - 2,200ppm threshold that might trigger the cloud feedback. Liu and Raftery show the probability of different emissions scenarios:

³⁸⁵ “Average resolution for the interval from 0 to 34 Ma is one sample every 2 ky; for the interval from 34 to 67 Ma, it is one sample every 4.4 kyr.” Thomas Westerhold et al., ‘An Astronomically Dated Record of Earth’s Climate and Its Predictability over the Last 66 Million Years’, *Science*, 11 September 2020, <https://www.science.org/doi/abs/10.1126/science.aba6853>.

³⁸⁶ “History has seen a multitude of studies (e.g., Svensmark, 1998; Lindzen et al., 2001; Schwartz, 2007) mostly implying lower ECS than the range assessed as very likely here. However, there are also examples of the opposite such as very large ECS estimates based on the Pleistocene records (Snyder, 2016), which has been shown to be overestimated due to a lack of accounting for orbital forcing and long term ice sheet feedbacks (Schmidt et al., 2017b), or suggestions that global climate instabilities may occur in the future (Steffen et al., 2018; Schneider et al., 2019). There is, however, no evidence for unforced instabilities of such magnitude occurring in the paleo record temperatures of the past 65 million years (Westerhold et al., 2020), possibly short of the PETM excursion (Chapter 5, Section 5.3.1.1) that occurred at more than 10°C above present (Anagnostou et al., 2020). Looking back, the resulting debates have led to a deeper understanding, strengthened the consensus, and have been scientifically valuable.” IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 7.5.5.

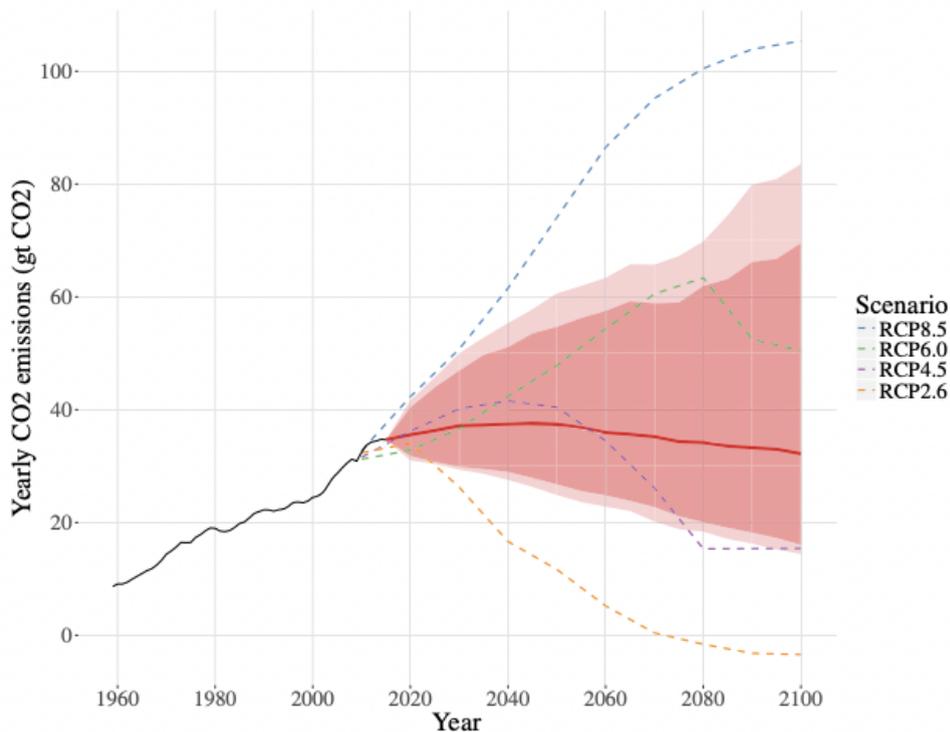
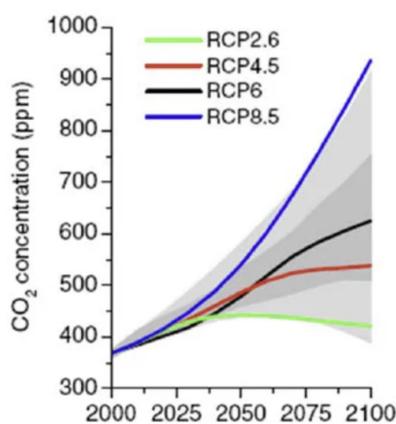


Fig. 1 Updated probabilistic forecast of CO₂ Emissions, based on data to 2015 and the method of Raftery et al.¹ The forecast median of yearly global emissions in 2100 is now 34 Giga tons.

Peiran R. Liu and Adrian E. Raftery, '[Country-Based Rate of Emissions Reductions Should Increase by 80% beyond Nationally Determined Contributions to Meet the 2 °C Target](#)', *Nature Communications Earth & Environment* 2, no. 1 (9 February 2021): Figure 1.

CO₂ concentrations on these different pathways are shown below



Thus, the probability of getting to 1,200ppm before 2100 now seems extremely small, probably below 0.1%. We could pass the threshold if we burned all the fossil fuels, on some estimates of remaining fossil fuels, but this would take several centuries, and seems extremely unlikely. As I discuss in Chapter 1, I put the probability at 1 in 500,000.

8.8.2. What effects would rapid cloud feedbacks have on human society?

It is very hard to say what effect the rapid cloud feedback would have on human society; this question has not been studied and it is very far outside the sample of human experience. Nevertheless, we have to try and form our best judgements with the information we have.

One way to guide thinking about this is to consider evidence from the paleoclimate. Much less intelligent and numerous species made it through when temperatures were 17°C higher during the Paleocene-Eocene Thermal Maximum, so this might make us optimistic about our own prospects. However, there are many disanalogies between our world and the Paleocene-Eocene Thermal Maximum, the most important being that we are reliant on agriculture. So: what effect would the cloud tipping point have on agriculture?

The major food crops would pass lethal limits in major food producing regions in the tropics and subtropics, making agriculture at least very difficult in regions that currently support billions of people.³⁸⁷ However, there would perhaps be scope to mitigate some of these effects by switching to heat-tolerant crops. Crops would not pass lethal limits in all regions, and land would be freed up at higher latitudes in Canada and Russia. So, agriculture would probably not stop entirely at the global level.

Droughts would likely become much worse across much of the world. This would have dire humanitarian consequences in certain regions. However, the general pattern is that climate change makes wet areas wetter and dry areas dryer. Moreover, many regions would be able to make use of irrigation to adapt to decreasing soil moisture. So, the effects on drought seem very unlikely to make agriculture impossible.

The final important factor is heat stress. Warming of 12°C would make the tropics and subtropics essentially uninhabitable due to heat stress and would make agriculture impossible (see Chapters 5 and 6). I am not aware of any studies of the effects of 13°C of heat stress on labour productivity. However, if we extrapolate the chart below, this suggests that for 13°C of warming, outdoor labour would be all but impossible in the tropics, and would be reduced by around 50% in temperate regions.

³⁸⁷ See Chapter 5.

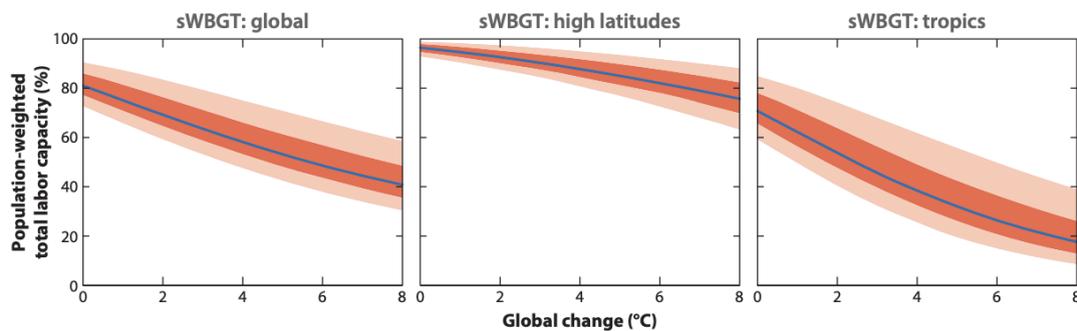


Figure 10

Population weighted total labor capacity. The CMIP5 ensemble is represented by the median (*blue line*), 50% (*red swath*), and 80% (*pink swath*) confidence intervals. The relative impacts on labor are shown at global (57°S to 57°N), high latitude (outside of 30°S to 30°N), and tropic (30°S to 30°N) regions. Abbreviations: CMIP, Coupled Model Intercomparison Project; sWBGT, simplified wet bulb globe temperature.

Source: Buzan and Huber, 'Moist Heat Stress on a Hotter Earth', 2020.

One possibly important mitigating factor is that agriculture may improve over the next few centuries in ways that we do not yet understand. I discuss the prospects for agriculture under extreme warming scenarios in Chapter 5.

Overall, the rapid cloud feedback would have extremely bad humanitarian effects, especially on the tropics, and my best guess is that it would lead to the starvation of billions of millions of people. Although this question has not been studied in depth and I am very uncertain, my best guess is that agriculture would still be possible and civilisation would survive, albeit greatly diminished.

Moreover, people would be stuck with the extreme greenhouse world for tens of thousands of years. So, the risk of long-run stagnation would increase, which would extend the so-called 'time of perils': the period in which we have the technology to destroy ourselves, but lack the political institutions necessary to manage that technology. This would also make recovery from catastrophe (from e.g. nuclear war or engineered pandemic) much harder.

Moreover, with such extreme warming, the indirect risks, such as mass migration and conflict would also be much higher. I discuss indirect risks in Chapters 11-13.

8.9. Hothouse Earth

Steffen et al's 'Trajectories of the Earth System in the Anthropocene' (2018) has been cited nearly 2,000 times at the time of writing. It argues that 2°C of warming is a planetary boundary beyond which we could enter into a 'Hothouse Earth' driven chiefly by nonlinear climate tipping points, rather than by fossil fuel burning.³⁸⁸

³⁸⁸ Will Steffen et al., 'Trajectories of the Earth System in the Anthropocene', *Proceedings of the National Academy of Sciences* 115, no. 33 (14 August 2018): 8252–59, <https://doi.org/10.1073/pnas.1810141115>.

The central claim of the paper is unclear because most of its central claims are about the tipping points that ‘might’ or ‘could’ be triggered due to global warming,³⁸⁹ or about the fact there might be a risk of such a tipping point.³⁹⁰ However, the claim “x could happen” is extremely broad, and the fact that we cannot exclude a risk doesn’t tell us anything about the magnitude of the risk.

It is also not clear what ‘Hothouse Earth’ means. They provide the following diagram which shows the purported 2°C tipping point.

³⁸⁹ “However, here we suggest that biogeophysical feedback processes within the Earth System coupled with direct human degradation of the biosphere may play a more important role than normally assumed, limiting the range of potential future trajectories and potentially eliminating the possibility of the intermediate trajectories” 8253; “Beyond this threshold, intrinsic biogeophysical feedbacks in the Earth System (Biogeophysical Feedbacks) could become the dominant processes controlling the system’s trajectory” Steffen et al., ‘Trajectories of the Earth System in the Anthropocene’, p. 8254.

³⁹⁰ “This analysis implies that, even if the Paris Accord target of a 1.5 °C to 2.0 °C rise in temperature is met, we cannot exclude the risk that a cascade of feedbacks could push the Earth System irreversibly onto a “Hothouse Earth” pathway”, Steffen et al., ‘Trajectories of the Earth System in the Anthropocene’, p. 8254.

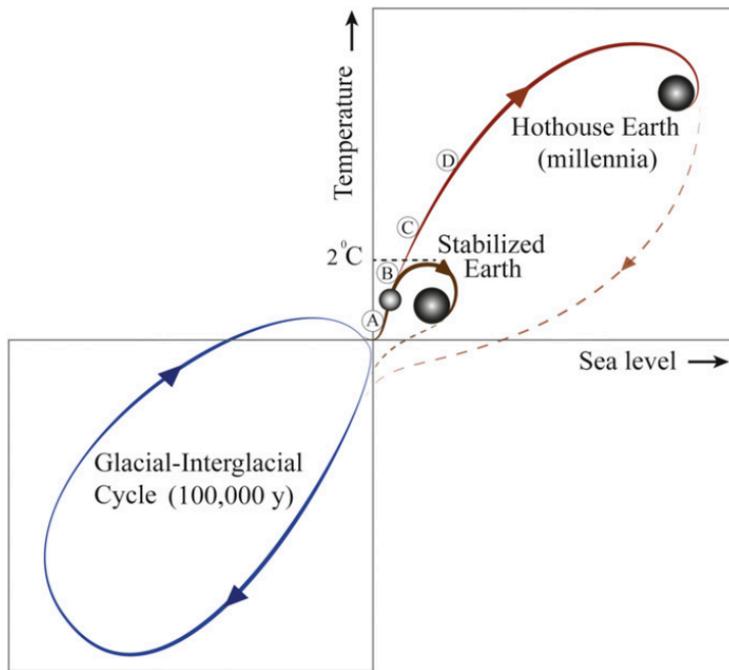


Fig. 1. A schematic illustration of possible future pathways of the climate against the background of the typical glacial-interglacial cycles (Lower Left). The interglacial state of the Earth System is at the top of the glacial-interglacial cycle, while the glacial state is at the bottom. Sea level follows temperature change relatively slowly through thermal expansion and the melting of glaciers and ice caps. The horizontal line in the middle of the figure represents the preindustrial temperature level, and the current position of the Earth System is shown by the small sphere on the red line close to the divergence between the Stabilized Earth and Hothouse Earth pathways. The proposed planetary threshold at $\sim 2^\circ\text{C}$ above the preindustrial level is also shown. The letters along the Stabilized Earth/Hothouse Earth pathways represent four time periods in Earth's recent past that may give insights into positions along these pathways (SI Appendix): A, Mid-Holocene; B, Eemian; C, Mid-Pliocene; and D, Mid-Miocene. Their positions on the pathway are approximate only. Their temperature ranges relative to preindustrial are given in SI Appendix, Table S1.

This diagram suggests that the Hothouse Earth is a world that is 6°C above pre-industrial, and that this state would be reached over the course of millennia. They argue that Hothouse Earth ultimately calls into question “the habitability of the planet for humans”.³⁹¹ This is a strong claim and they do not argue for it.

Moreover, Steffen et al (2018) do not present any evidence or argument showing that there is a non-negligible risk that at 2°C , tipping points alone could take us up to 6°C over the course of millennia. The table below shows the combined effect of the feedbacks they mention in the main text:

³⁹¹ Steffen et al., ‘Trajectories of the Earth System in the Anthropocene’, p. 8256.

Table 1. Carbon cycle feedbacks in the Earth System that could accelerate global warming

Feedback	Strength of feedback by 2100,* °C	Refs. (<i>SI Appendix, Table S2</i> has more details)
Permafrost thawing	0.09 (0.04–0.16)	20–23
Relative weakening of land and ocean physiological C sinks	0.25 (0.13–0.37)	24
Increased bacterial respiration in the ocean	0.02	25, 26
Amazon forest dieback	0.05 (0.03–0.11)	27
Boreal forest dieback	0.06 (0.02–0.10)	28
Total	0.47 (0.24–0.66)	

The strength of the feedback is estimated at 2100 for an ~2 °C warming.

*The additional temperature rise (degrees Celsius) by 2100 arising from the feedback.

By their own calculations, they find that 2°C of warming would cause an extra 0.47°C by 2100 from these feedbacks. They do not explain what might cause the remaining 3.5°C of warming over millennia implied by their Figure 1. Furthermore, some of the numbers in their table are not supported by the latest evidence:

- The weakening of land and ocean carbon sinks is a feedback that is accounted for in IPCC estimates of the transient climate response to cumulative emissions. Although carbon sinks weaken as temperature rises, CO₂ concentrations have a diminishing effect on warming as they increase. The evidence suggests that these effects cancel out such that temperatures increase in proportion to cumulative emissions, and when emissions stop, the evidence suggests that temperatures will be roughly constant for 100 years before slowly declining over millennia.³⁹² Thus, the consensus in the literature is that warming past 2°C would not cause additional warming, unless there were further CO₂ emissions.
- As discussed above, the IPCC Sixth Assessment Report suggests that boreal forest dieback will have roughly neutral net impact on emissions.
- Regarding Amazon forest dieback, most models now suggest that the effect of CO₂ fertilisation overwhelms the effect of warming and precipitation change with respect to changes in the Amazon biome. On the very pessimistic assumption of no CO₂ fertilisation, the IPCC finds that for 2°C, Amazon forest dieback would increase emissions by up to 100GtC, which would increase temperatures by 0.017°C,³⁹³ which is below outside the range suggested in Steffen et al's Table 1.

Again, there is evidence from the paleoclimate record that once we pass 2°C, feedback effects alone will not cause further warming of 4°C. During the Pliocene, temperatures were 3°C warmer than pre-industrial without tipping the climate system into a Hothouse Earth state.

³⁹² “The Zero Emissions Commitment (ZEC) is the change in global mean temperature expected to occur following the cessation of net CO₂ emissions and as such is a critical parameter for calculating the remaining carbon budget... Overall, the most likely value of ZEC on multi-decadal timescales is close to zero, consistent with previous model experiments and simple theory.” Andrew H. MacDougall et al., ‘Is There Warming in the Pipeline? A Multi-Model Analysis of the Zero Emissions Commitment from CO₂’, *Biogeosciences* 17, no. 11 (15 June 2020): Figure 3b. <https://doi.org/10.5194/bg-17-2987-2020>.

³⁹³ This is using the IPCC's estimate of the transient climate response to cumulative emissions, which is 1.65°C per 1,000 billion tonnes of carbon.

If we look at the specific feedbacks considered by Steffen et al (2018), there is no evidence for these dramatic tipping points in our climate past: the rainforest flourished during the Eocene and there was no abrupt release of permafrost carbon during the Pliocene when the poles were 10°C warmer than today.

The IPCC also suggests that there is little evidence in the paleoclimate of the extreme climate sensitivity values implied by Steffen et al (2018).

“History has seen a multitude of studies (e.g., Svensmark, 1998; Lindzen et al., 2001; Schwartz, 2007) mostly implying lower ECS than the range assessed as very likely here. However, there are also examples of the opposite such as very large ECS estimates based on the Pleistocene records (Snyder, 2016), which has been shown to be overestimated due to a lack of accounting for orbital forcing and long term ice sheet feedbacks (Schmidt et al., 2017b), or suggestions that global climate instabilities may occur in the future (Steffen et al., 2018; Schneider et al., 2019). There is, however, no evidence for unforced instabilities of such magnitude occurring in the paleo record temperatures of the past 65 million years (Westerhold et al., 2020), possibly short of the PETM excursion (Chapter 5, Section 5.3.1.1) that occurred at more than 10°C above present (Anagnostou et al., 2020). Looking back, the resulting debates have led to a deeper understanding, strengthened the consensus, and have been scientifically valuable.”³⁹⁴

³⁹⁴ IPCC, *Climate Change 2021: The Physical Science Basis*, Sixth Assessment Report (UNFCCC, 2021), sec. 7.5.5.

9. Greenhouses and habitability

Some worry that climate change could directly bring about human extinction. The main routes to extinction that have been discussed in the literature are the ‘runaway greenhouse effect’ and the risk that Earth transitions into a ‘moist greenhouse’.

In *The Precipice*, Toby Ord estimates that the probability that climate change could directly cause extinction in the next 100 years is 1 in 1,000.³⁹⁵ This is largely on the basis of the possibility that we could, by burning fossil fuels, cause a moist greenhouse or runaway greenhouse. I will argue that the risk of human extinction due to the loss of water on Earth is several orders of magnitude lower than this. To be clear, my argument is not (only) about the probability of extinction given a certain amount of emissions, but is rather about the all-things-considered probability of extinction, taking into account the probability of different emissions scenarios.

Planetary sterilisation by the loss of water to space is extremely unlikely to kill all humans. Other mechanisms, such as the destruction of agriculture are more probable, but still very unlikely.

Disclosure: I carried out some research for the climate change section of *The Precipice*.

9.1. The long-term habitability of Earth

The Sun pours tremendous amounts of energy on to the Earth. Around a third is reflected back to space by the atmosphere, while the remainder hits the Earth surface and becomes heat. Greenhouse gases are more transparent to the shortwave radiation coming from the Sun than they are to the longwave radiation (heat) emitted by the surface of the earth. Consequently, greenhouse gases like water vapour, CO₂ and methane act like a quilt, trapping heat and warming the planet. The greenhouse effect makes the Earth 33°C warmer than it would otherwise be and so makes life on Earth possible. This same physical process will eventually make the planet uninhabitable.

The Sun is increasing in brightness by 1% every 110 million years.³⁹⁶

“In 7.6 billion years, the Sun will have grown so vast that it will balloon out beyond Earth’s own orbit, either swallowing our planet or flinging it out much further. And either way, in 8 billion years our Sun itself will die.”³⁹⁷

Well before then, the Earth will become uninhabitable.

Warming climates are thought to transition through two distinct phases: the *moist greenhouse* and the *runaway greenhouse*. A ‘moist greenhouse’ refers to hot water-rich atmospheres, in which all of the Earth’s water is eventually lost to space. The standard

³⁹⁵ Toby Ord, *The Precipice: Existential Risk and the Future of Humanity* (Bloomsbury Publishing, 2020), p. 167.

³⁹⁶ E. T. Wolf and O. B. Toon, ‘Delayed Onset of Runaway and Moist Greenhouse Climates for Earth’, *Geophysical Research Letters* 41, no. 1 (2014): 167–72, <https://doi.org/10.1002/2013GL058376>.

³⁹⁷ Ord, *The Precipice*, Ch. 8.

threshold for a moist greenhouse is a global mean surface temperature of 67°C (compared to 15°C today).³⁹⁸

A runaway greenhouse differs from a moist greenhouse. In a runaway greenhouse, all of the Earth's water is contained in the atmosphere as vapour and cloud. Due to the thickening quilt of water vapour in the atmosphere, the Earth is unable to cool down through releasing long-wave radiation to space. Thus, the planet warms uncontrollably until all the surface water has evaporated and surface temperatures reach 1,300°C.³⁹⁹

A runaway greenhouse probably happened on Venus. Venus' atmosphere is [almost pure CO₂](#). Were it not for the greenhouse effect, the surface temperature of Venus would be -42°C. In reality, the surface of Venus is about 470°C - [hot enough to melt lead](#). This is even hotter than Mercury (a relatively temperate 166°C) - even though Mercury is closer to the Sun than Venus.

A moist greenhouse is generally believed to precede a runaway greenhouse and so places a more stringent limit on planetary habitability. Simulations with modern models suggest that we are safe from a moist greenhouse for at least a billion years, though more simple models suggest that we could reach a moist greenhouse in hundreds of millions of years.⁴⁰⁰ Once we reached a moist greenhouse threshold, it would take hundreds of millions of years to lose the oceans to space.

9.2. Could humans cause a moist greenhouse or runaway greenhouse?

I argued in Chapter 1 that ultimately recoverable fossil fuels are 1-3 trillion tonnes of carbon. If we were to burn 3 trillion tonnes of carbon, CO₂ concentrations would increase to around 1,600ppm.⁴⁰¹ This is markedly lower than estimates of recoverable fossil fuels sometimes used in climate science. For example, some studies model the effects of burning more than

³⁹⁸ “One-dimensional radiative-convective calculations predict that stratospheric water vapor mixing ratios of $\sim 3 \times 10^{-3}$ can occur if the global mean surface temperature reaches 340 K, under the assumptions of water vapor saturation and a 200 K isothermal stratosphere [Kasting, 1988]. This description is often taken as the standard threshold for a moist greenhouse climate. A moist greenhouse is generally believed to precede a runaway greenhouse in the evolutionary sequence of warming terrestrial atmospheres. While a moist greenhouse is climatologically stable, unlike a runaway greenhouse, the potential for rapid water loss makes it the more proximal boundary to the inner edge of the habitable zone [Kasting et al., 1993; Kopparapu et al., 2013].” (wolf and toon 2015)
³⁹⁹ E. T. Wolf and O. B. Toon, ‘The Evolution of Habitable Climates under the Brightening Sun’, *Journal of Geophysical Research: Atmospheres* 120, no. 12 (2015): 5775–94, <https://doi.org/10.1002/2015JD023302>.

⁴⁰⁰ “Nonetheless, our results imply that Earth's climate may remain safe against both water loss and thermal runaway limits for at least another 1.5 billion years and probably for much longer.” “Cloud-free 1-D models with saturated atmospheres predict that Earth will reach moist greenhouse conditions (Ts = 340 K) when the solar constant increases by only 1.5% above its present level [Kopparapu et al., 2013]. Thus, our home planet may be subject to a moist greenhouse climate in a mere ~170 million years [Gough, 1981].” Wolf and Toon, ‘Delayed Onset of Runaway and Moist Greenhouse Climates for Earth’.

⁴⁰¹ N. S. Lord et al., ‘An Impulse Response Function for the “Long Tail” of Excess Atmospheric CO₂ in an Earth System Model’, *Global Biogeochemical Cycles* 30, no. 1 (2016): 2–17, <https://doi.org/10.1002/2014GB005074>.

10 trillion tonnes of carbon,⁴⁰² which would increase CO₂ concentrations past 5,000ppm.⁴⁰³ The lower estimate of recoverable fossil fuels has important implications for the likelihood of moist or runaway greenhouse on Earth.

The table below shows the findings of different models on the conditions required to trigger a runaway greenhouse.

Study	Findings
Kasting and Ackerman (1986) ⁴⁰⁴	Slightly over 10,000ppm would be required to cause a moist greenhouse.
Goldblatt et al (2013) ⁴⁰⁵	Anthropogenic emissions could “in theory” produce runaway greenhouse, but are probably insufficient. 30,000ppm may trigger a runaway.
Hansen et al (2013) ⁴⁰⁶	No plausible human-made greenhouse gas forcing can cause a runaway greenhouse effect. Warming of more than 16-24°C could produce a moist greenhouse, but natural weathering would remove the excess atmospheric CO ₂ on a time scale of 1,000 to 10,000 years, well before the ocean is significantly depleted.
Ramirez et al (2014) ⁴⁰⁷	On “the most alarmist assumptions possible”, the model nearly runs away at 3,300ppm. But on more plausible assumptions, runaway appears to be impossible from fossil fuel emissions.
Wolf and Toon (2014) ⁴⁰⁸	A 15.5% increase in solar forcing doesn’t lead to moist greenhouse or runaway greenhouse. Since each 2% increase in solar forcing is equivalent to a doubling of CO ₂ , this is equivalent to nearly 8 doublings, which is CO ₂ concentrations of 70,000ppm.
Wolf and Toon (2015) ⁴⁰⁹	Significant water loss begins to happen if solar forcing increases by 19%. This is equivalent to more than nine doublings of pre-industrial CO ₂ , or CO ₂ concentrations of 140,000ppm.

⁴⁰² Ricarda Winkelmann et al., ‘Combustion of Available Fossil Fuel Resources Sufficient to Eliminate the Antarctic Ice Sheet’, *Science Advances* 1, no. 8 (1 September 2015): e1500589, <https://doi.org/10.1126/sciadv.1500589>.

⁴⁰³ N. S. Lord et al., ‘An Impulse Response Function for the “Long Tail” of Excess Atmospheric CO₂ in an Earth System Model’, *Global Biogeochemical Cycles* 30, no. 1 (2016): 2–17, <https://doi.org/10.1002/2014GB005074>.

⁴⁰⁴ J. F. Kasting and T. P. Ackerman, ‘Climatic Consequences of Very High Carbon Dioxide Levels in the Earth’s Early Atmosphere’, *Science* 234, no. 4782 (12 December 1986): 1383–85, <https://doi.org/10.1126/science.11539665>.

⁴⁰⁵ Colin Goldblatt et al., ‘Low Simulated Radiation Limit for Runaway Greenhouse Climates’, *Nature Geoscience* 6, no. 8 (August 2013): 661–67, <https://doi.org/10.1038/ngeo1892>

⁴⁰⁶ James Hansen et al., ‘Climate Sensitivity, Sea Level and Atmospheric Carbon Dioxide’, *Phil. Trans. R. Soc. A* 371, no. 2001 (28 October 2013): 20120294, <https://doi.org/10.1098/rsta.2012.0294>.

⁴⁰⁷ Ramses M. Ramirez et al., ‘Can Increased Atmospheric CO₂ Levels Trigger a Runaway Greenhouse?’, *Astrobiology* 14, no. 8 (1 August 2014): 714–31, <https://doi.org/10.1089/ast.2014.1153>

⁴⁰⁸ Wolf and Toon, ‘Delayed Onset of Runaway and Moist Greenhouse Climates for Earth’.

⁴⁰⁹ E. T. Wolf and O. B. Toon, ‘The Evolution of Habitable Climates under the Brightening Sun’, *Journal of Geophysical Research: Atmospheres* 120, no. 12 (2015): 5775–94, <https://doi.org/10.1002/2015JD023302>.

Popp et al (2016) ⁴¹⁰	Increasing concentrations to 1,500ppm causes a transition to a moist greenhouse. However, this assumes that the world is made entirely of ocean that is only 50 metres deep. Adjusting for this, the forcing is more like 4,400ppm. Moreover, Max Popp told me in personal correspondence that their paper is more relevant to the habitability of exoplanets than current or near-future climate change, and that their paper doesn't suggest that CO2 concentrations of 4,000ppm would trigger a transition into a moist greenhouse climate. ⁴¹¹
----------------------------------	---

In his book *Six Degrees*, Mark Lynas argues that six degrees of warming could possibly result in human extinction due to the risk of runaway feedbacks leading to a moist greenhouse or runaway greenhouse. So, his claim is not that the impacts of 6°C of warming would themselves directly cause human extinction, but rather 6°C of warming might cause further warming which could then cause human extinction. Specifically, he mentions the cloud feedback found in Schneider et al (2019), according to which once CO₂ concentrations pass 1,200ppm, there is warming of 8°C on top of the warming we have already experienced, which would be around 5°C. On top of this, other feedbacks from permafrost, methane clathrates or forest dieback could kick in. All told, this could bring us perilously close to a runaway greenhouse.

As we have seen, this is strongly at odds with the findings from models. For example, Lynas reports the finding of the Ramirez et al (2014) study without also pointing out that they only get a runaway at 3,300ppm on “the most alarmist assumptions possible”.⁴¹²

Moreover, there is strong reason from the paleoclimate record to believe that, even if the finding of the Schneider et al (2019) paper is correct, the risk of triggering a runaway by pushing warming to 6°C is minimal. If this feedback were real, then it would have kicked in during the Paleocene-Eocene Thermal Maximum when CO₂ concentrations were above 2,000ppm and temperatures were 17°C above pre-industrial levels, or the early Eocene when temperatures were 10°C than pre-industrial levels, or during the mid-Cretaceous when temperatures were 20°C higher, or during the end-Permian when temperatures were upwards of 17°C higher. The other feedbacks Lynas mentions would also have kicked in at these times: since there were no ice caps in all of these periods, if huge amounts of permafrost carbon and methane clathrates are going to be released in the future, they would have been released in these periods. In none of these cases did we trigger feedback processes that caused a feedback that killed all life on Earth. Therefore, the paleoclimate evidence agrees with the models that such a scenario is extremely unlikely.

One difference mentioned by Lynas is that the Sun is getting hotter by 1% every 110 million years. So, since the Eocene, the Sun would be 0.5% hotter, which roughly translates to half

⁴¹⁰ Max Popp, Hauke Schmidt, and Jochem Marotzke, ‘Transition to a Moist Greenhouse with CO₂ and Solar Forcing’, *Nature Communications* 7, no. 1 (9 February 2016): 10627, <https://doi.org/10.1038/ncomms10627>.

⁴¹¹ Max Popp, personal correspondence, 11th and 12th April 2021.

⁴¹² Mark Lynas, *Our Final Warning: Six Degrees of Climate Emergency* (London: 4th Estate, 2020), 266–67.

a degree of additional warming.⁴¹³ This would only make a small difference to the likelihood of a runaway or moist greenhouse.

There might also be other differences that mean that in previous super-greenhouse periods, the planet was lucky to avoid a runaway.⁴¹⁴ For example, the Earth's orbit in relation to the Sun or the Earth's position on its axis might have changed, which places us at greater risk today. It is difficult to completely eliminate residual uncertainty about a potential runaway. Goldblatt and Watson (2012) provide a nice summary of the argument from residual uncertainty:

“Here, we review what is known about the runaway greenhouse to answer this question, describing the various limits on outgoing radiation and how climate will evolve between these. The good news is that almost all lines of evidence lead us to believe that is unlikely to be possible, even in principle, to trigger full a runaway greenhouse by addition of non-condensable greenhouse gases such as carbon dioxide to the atmosphere. However, our understanding of the dynamics, thermodynamics, radiative transfer and cloud physics of hot and steamy atmospheres is weak. We cannot therefore completely rule out the possibility that human actions might cause a transition, if not to full runaway, then at least to a much warmer climate state than the present one”⁴¹⁵

I am not sure how to quantify and bound this kind of extreme residual model uncertainty.

I spoke to one researcher about this, and they estimated that if we burn all the fossil fuels, the chance of a moist or runaway greenhouse is around 10%. This seems too high to me given the evidence from the paleoclimate and the models, but it makes sense for me to defer to them on this. It would be useful to have expert elicitation studies on this question.

⁴¹³ This is calculated as follows. A 2% increase in solar irradiance is equivalent to one doubling of CO₂. Hansen et al., 'Climate Sensitivity, Sea Level and Atmospheric Carbon Dioxide', p. 2.

A doubling of CO₂ is equivalent to an increase in radiative forcing of around 3.7 watts per square metre. A. Farnsworth et al., 'Climate Sensitivity on Geological Timescales Controlled by Nonlinear Feedbacks and Ocean Circulation', *Geophysical Research Letters* 46, no. 16 (2019): 9880–89, <https://doi.org/10.1029/2019GL083574>.

So, increasing solar irradiance by 0.5% would be equivalent to increasing radiative forcing by around 0.9 watts per square metre. There is around 0.7 degrees of warming for every additional watt per square metre of radiative forcing. Jeffrey T. Kiehl and Christine A. Shields, 'Sensitivity of the Palaeocene–Eocene Thermal Maximum Climate to Cloud Properties', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371, no. 2001 (28 October 2013): 20130093, <https://doi.org/10.1098/rsta.2013.0093>.

So, the additional warming from the brightening Sun since the early Eocene would cause an additional 0.6 degrees of warming.

⁴¹⁴ Thanks to Andrew Watson for discussion of this point.

⁴¹⁵ Colin Goldblatt and Andrew J. Watson, 'The Runaway Greenhouse: Implications for Future Climate Change, Geoengineering and Planetary Atmospheres', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370, no. 1974 (13 September 2012): 4197–4216, <https://doi.org/10.1098/rsta.2012.0004>.

One counter-argument to this is that this argument from extreme model uncertainty also applies in the other direction. There is also some risk that in the absence of CO₂ emissions, the world will enter a glacial period. This could also threaten civilisation because it would make agriculture very difficult. Our models and paleoclimate evidence suggest that given how high CO₂ concentrations now are, re-entering a glacial period is extremely unlikely.⁴¹⁶ But again, there is some residual model uncertainty that is difficult to eliminate. It seems plausible to me that these kinds of model uncertainty largely cancel out; it is not clear that the risk is greater in one direction rather than the other.

9.3. How might a moist or runaway greenhouse kill us?

If we did produce a moist or runaway greenhouse on Earth, how would it kill us? One possibility is that the planet would be sterilised because the Earth's water would be lost to space. However, this kill mechanism seems unlikely to be the one that destroys civilisation.

Even if we did transition to a moist greenhouse, it would take hundreds of millions of years to lose the Earth's oceans to space. However, once man-made emissions stop, natural sequestration processes would reduce CO₂ concentrations close to natural levels after tens of thousands of years. CO₂ concentrations would look like this in the long-term:

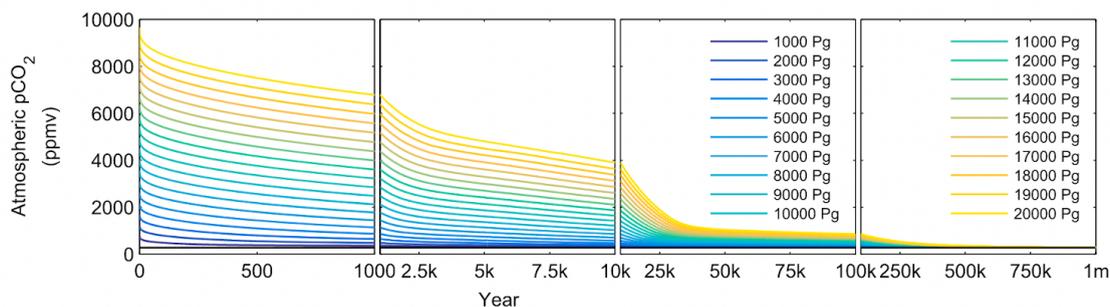


Figure 2. Atmospheric pCO₂ predicted by cGENIE for the pulse series scenarios (1000–20,000 Pg C). Preindustrial CO₂ concentrations are shown in black.

Source: N. S. Lord et al., 'An Impulse Response Function for the "Long Tail" of Excess Atmospheric CO₂ in an Earth System Model', *Global Biogeochemical Cycles* 30, no. 1 (2016): Fig. 2, <https://doi.org/10.1002/2014GB005074>.

The Earth's water would be lost to space only if CO₂ concentrations remain high for millions of years, but in fact they will fall back close to natural levels much sooner.⁴¹⁷ Even if we were

⁴¹⁶ "It is virtually certain that orbital forcing will be unable to trigger widespread glaciation during the next 1000 years. Paleoclimate records indicate that, for orbital configurations close to the present one, glacial inception only occurred for atmospheric CO₂ concentrations significantly lower than pre-industrial levels. Climate models simulate no glacial inception during the next 50,000 years if CO₂ concentrations remain above 300 ppm." IPCC, *Climate Change 2013: The Physical Science Basis*, Fifth Assessment Report, p. 399.

⁴¹⁷ "For simulations +19%, +20%, and +21% S₀, the entirety of Earth's oceans could evaporate in as little as ~3.5 Gyr, ~672 Myr, and ~130 Myr given a static climate. However, more detailed calculations of atmospheric water loss from moist greenhouse atmospheres are warranted." E. T. Wolf and O. B. Toon, 'The Evolution of Habitable Climates under the Brightening Sun', *Journal of Geophysical Research: Atmospheres* 120, no. 12 (2015): 5775–94, <https://doi.org/10.1002/2015JD023302>.

to burn through all of the fossil fuels, we would not be killed off by the oceans being lost to space.

Even so, the moist or runaway greenhouse would make the planet uninhabitable long before the Earth's oceans are lost to space. Recall that the threshold for a moist greenhouse is surface temperatures of around 66°C, compared to 15°C today. If the world were this hot, it is very difficult to see how human civilisation could survive - crops would pass lethal limits almost everywhere and unlivable heat stress would envelop the world.

9.4. Quantifying the direct extinction risk of climate change

It is very difficult to estimate the overall probability that climate change would directly cause human extinction. It is important to note that Toby Ord's '1 in 1,000' estimate is for the next 100 years. But if we are going to cause a moist or runaway greenhouse, we would have to burn huge amounts of carbon, which would most likely take several centuries. This is one reason that the 1 in 1,000 figure is probably too high.

I will explore two ways of quantifying the direct extinction risk.

1. **Runaway feedbacks:** Assuming that moist or runaway greenhouse feedback effects would almost certainly kill life on Earth, what is the probability that we trigger such feedbacks?
2. **Extinction-level warming thresholds:** Define some warming thresholds past which humans would go extinct, and calculate the probability of passing such thresholds.

Runaway feedbacks

First consider the approach of calculating the risk of triggering runaway feedbacks. We saw above that one expert thought if we burn all the fossil fuels (releasing 3,000 GtC and causing concentrations to rise to 1,600ppm), then the risk of a runaway or moist greenhouse is 10%. This seems too high to me given model findings and paleoclimate evidence. But assuming this is true, what is the probability that we could trigger these extreme feedbacks?

I argued in section 1.5 that it is difficult to come up with plausible scenarios in which we burn 3,000 GtC because it would require dramatic progress in advanced coal extraction

"High stratospheric H₂O would also result in increased rates of water loss by way of photodissociation followed by hydrogen escape—a phenomenon sometimes termed a moist greenhouse (Kasting, 1988). However, atmospheric CO₂ concentrations would presumably be restored to more normal values by silicate weathering within a few million years (Walker et al., 1981), before substantial water loss could occur." Ramses M. Ramirez et al., 'Can Increased Atmospheric CO₂ Levels Trigger a Runaway Greenhouse?', *Astrobiology* 14, no. 8 (1 August 2014): 714–31, <https://doi.org/10.1089/ast.2014.1153>.

"A warming of 16–24°C produces a moderately moist greenhouse, with water vapour increasing to about 1% of the atmosphere's mass, thus increasing the rate of hydrogen escape to space. However, if the forcing is by fossil fuel CO₂, the weathering process would remove the excess atmospheric CO₂ on a time scale of 10⁴–10⁵ years, well before the ocean is significantly depleted. Baked-crust hothouse conditions on the Earth require a large long-term forcing that is unlikely to occur until the sun brightens by a few tens of per cent, which will take a few billion years [121]." James Hansen et al., 'Climate Sensitivity, Sea Level and Atmospheric Carbon Dioxide', *Phil. Trans. R. Soc. A* 371, no. 2001 (28 October 2013): 20120294, <https://doi.org/10.1098/rsta.2012.0294>, p.24

technology and societal willingness to deploy the technology, but little progress in low carbon technologies. There are no estimates in the literature of such extreme emissions scenarios. Liu and Raftery (2021) estimate a 90% confidence interval for cumulative emissions to 2100 of roughly 550 GtC to 1,200 GtC.⁴¹⁸ Thus, it at least seems fair to say that the probability of 3,000GtC by 2100 is well below 1%.

My own rough back of the envelope [model](#) suggests that the chance over all time is 1 in 500,000. Combined with the probability of extinction conditional on burning all the fossil fuels, the risk of direct extinction from climate change is 1 in 5 million.

Extinction-level warming thresholds

Martin Weitzman defends using 10°C as an ‘illustrator threshold’ of a level of warming that would threaten human civilisation. Weitzman relies mainly on the paper on heat stress discussed by Sherwood and Huber (2010) discussed in Chapter 5, which found that at 12°C, the majority of the population would be exposed to lethal heat stress, as people are currently distributed.

“Thus, a temperature change of 10°C would appear to represent an extreme threat to human civilization and global ecology as we now know it, even if it might not necessarily mean the end of Homo sapiens as a species.”⁴¹⁹

I agree with Weitzman that it is difficult to see how 10°C could cause human extinction. The heat stress mechanism discussed by Sherwood and Huber (2010) would not itself cause human extinction since higher latitudes and altitudes would be spared, and there would be some scope for adaptation with air conditioning and migration.

It also seems that, despite rising heat stress, human labour would still be viable at higher latitudes, though it would likely be impossible in the tropics, as discussed in Chapter 8.

I discussed lethal limits to food production in Chapter 5. Research on the effects of more than 4.4°C of warming on agricultural production is limited. But crops would not pass lethal limits in the major food producing regions even if there were 10°C of warming. The level of warming required to completely destroy food production in the main food producing regions is in excess of 20°C.

If we pessimistically assume that the 10°C threshold is the correct threshold for extinction, it is difficult to quantify the risk of passing the threshold. The only published quantitative estimates in the literature project out to 2100 and find a much lower range of potential warming. The findings of Liu and Raftery (2021) are fairly representative of other studies in the literature and they find that the 90% confidence interval for warming by 2100 is 2°C to 3.9°C. The risk of 10°C of warming in the next 100 years seems extremely small.

⁴¹⁸ Liu and Raftery (2021) do not give explicit numbers in their paper, but this is inferred from their Figure 1. As I discussed in Chapter 1, I think this range is too pessimistic.

⁴¹⁹ Martin L. Weitzman, ‘Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change’, *Review of Environmental Economics and Policy* 5, no. 2 (7 January 2011): 275–92, <https://doi.org/10.1093/reep/rer006>.

My back of the envelope [model](#) suggests that the chance of burning all the fossil fuels over all time is 1 in 500,000. If we burn all of the recoverable fossil fuels (3,000 GtC), there is a 1 in 6 chance of more than 9.6°C. This suggests that over all time, the direct risk of extinction from climate change over all time is on the order of 1 in 3 million.

On the scenario in which we burn all of the fossil fuels, doing so takes around 400 years. The risk of burning all of the fossil fuels over the next 100 years is far smaller, plausibly by at least an order of magnitude. The 100-year direct extinction risk would be correspondingly smaller.

A related way to approach this might be to estimate the probability that we trigger the tipping point identified by Schneider et al (2019) that I discussed in Chapter 8. If Schneider et al (2019) is correct, then once concentrations pass 1,200ppm to 2,200ppm, a cloud feedback would cause warming of 8°C over decades. Suppose that this would cause extinction if it occurred.

Given the 'indefinitely stalled decarbonisation' model, the chance of passing 1,200ppm is plausibly on the order of 1 in 100,000. As discussed in Chapter 8, scientists are divided on how plausible the cloud feedback is. If we assume it has 50% probability, then the risk of extinction is around 1 in 200,000.

10. Economic costs

10.1. Context and trends

The context and trends for economic growth are discussed in Chapter 5. In brief:

- Over the course of the twentieth century, [world GDP](#) grew by 1,800%, and per capita world GDP rose by 448%.
- In one 2018 expert elicitation study, experts predicted that per capita GDP will rise by between 200% and 2,000% by 2100 (95% confidence interval).⁴²⁰
- However, some countries, predominantly in Sub-Saharan Africa, have essentially never experienced a sustained increase in living standards and billions of people there are still close to subsistence.

The range of GDP and GDP per capita on different Shared Socioeconomic Pathways is shown below

⁴²⁰ P. Christensen, K. Gillingham, and W. Nordhaus, 'Uncertainty in Forecasts of Long-Run Economic Growth.', *Proceedings of the National Academy of Sciences of the United States of America* 115, no. 21 (2018): fig. 1.

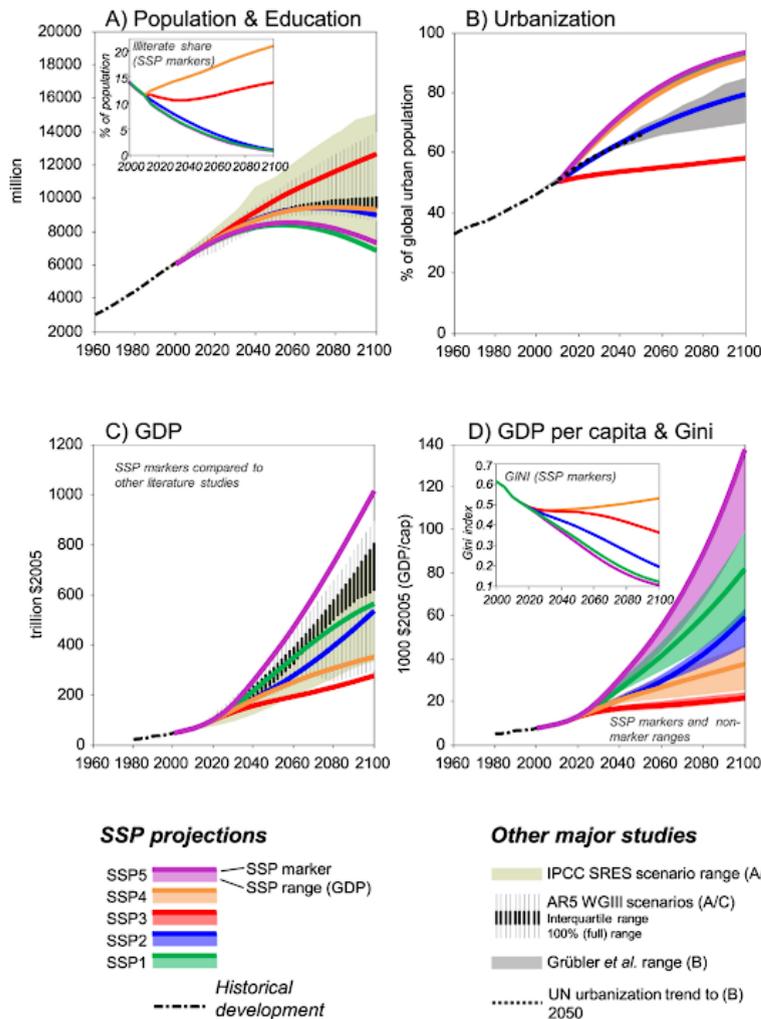


Fig. 2. Development of global population and education (A), urbanization (B), GDP (C), and GDP per capita and the Gini index (D). The inset in panel A gives the share of people without education at age of ≥ 15 years, and the inset in panel D denotes the development of the global (cross-national) Gini index. The SSPs are compared to ranges from other major studies in the literature, such as the IPCC AR5 (Clarke et al., 2014); IPCC SRES (Nakicenovic and Swart, 2000), UN, and Grubler et al. (2007). The colored areas for GDP (panel D) denote the range of alternative SSP GDP projections presented in this Special Issue (Dellink et al. (2016), Crespo Cuaresma (2016), Leimbach et al. (2016)).

As discussed in Chapter 5, given historical trends, the most plausible SSP is SSP4, though no scenarios do not account for the possibility of an acceleration of growth driven by advanced AI.

10.2. The social cost of carbon

The social cost of carbon is the net present value of future global climate change impacts from one additional tonne of CO₂ emitted to the atmosphere at a particular point in time. In short, it is the marginal cost of an additional tonne of CO₂. The social cost of carbon allows us to prioritise marginal climate change policy using cost-benefit analysis.

The social cost of carbon is computed by comparing two scenarios:

1. **Reference scenario:** projecting a future global socioeconomic scenario for centuries, and the resulting global greenhouse gas emissions, climate change, and net global damages.

2. **Pulsed scenario:** projecting a future scenario the same as the reference scenario, except with an additional pulse of CO₂ from some specific year, e.g. 2020.

The social cost of carbon is the discounted net present value of the additional net damages in the pulsed scenario compared to the reference scenario.⁴²¹

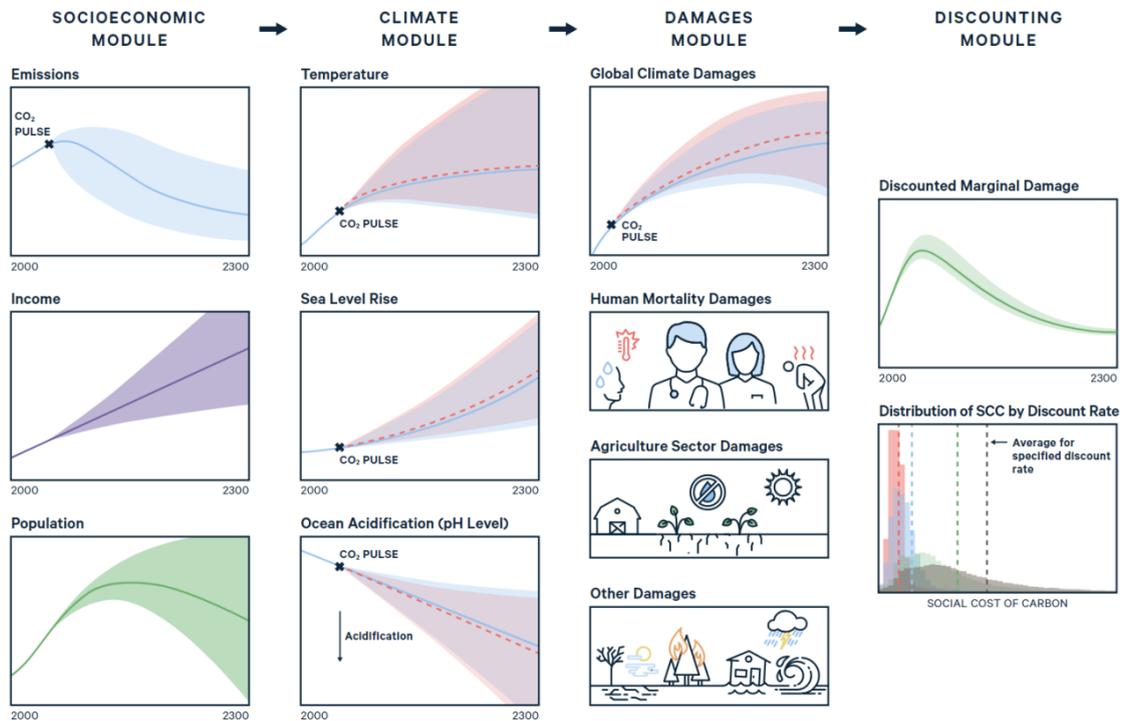
Models that assess the social cost of carbon are comprised of multiple different modules, including:

1. **Socio-economic module:** characterising the future evolution of the economy, including future CO₂ but without climate change.
2. **Climate module:** Quantifies how the earth system responds to emissions.
3. **Damages module:** Quantifies how the much damage changes in the Earth system do to the economy.
4. **Discounting module:** Converts the sum of future damages into a net present value. We may want to discount because people in the future are likely to be richer or because of an evaluative judgement that future people's welfare is worth less than that of present people.

This is illustrated in the following schematic

⁴²¹ "The SCC is conceptually the marginal cost to society of emitting carbon dioxide (CO₂). Computationally, the SCC is the net present value of future global climate change impacts from one additional net global metric ton of CO₂ emitted to the atmosphere at a particular point in time (Figure 1). An SCC value is computed using two long-run scenarios – a reference scenario projecting a future global socioeconomic condition for centuries, and the resulting global greenhouse gas emissions, climate change, and net global damages from that climate change; and, a pulsed scenario projecting the incremental climate change and damages over time from the addition of a pulse of CO₂ in an individual year (e.g., 2020) to the reference scenario. An SCC in 2020, therefore, is the discounted value of the additional net damages from the marginal emissions increase in 2020 relative to the reference condition." Steven K. Rose, Delavane B. Diaz, and Geoffrey J. Blanford, 'Understanding the Social Cost of Carbon: A Model Diagnostic and Inter-Comparison Study', *Climate Change Economics* 08, no. 02 (May 2017): 1750009, <https://doi.org/10.1142/S2010007817500099>.

Figure 2. Estimating Social Cost of Carbon under Uncertainty



Note: Estimation involves a baseline case (solid) versus a pulse of emissions (red dashed lines and areas). Shading depicts probability distributions on projections.

Source: Resources for the Future, *The Social Cost of Carbon: Advances in Long-Term Probabilistic Projections of Population, GDP, Emissions, and Discount Rates*, October 2021

My main focus here will be on the damage module, as it is more relevant to my main question and I have discussed the other modules elsewhere in this report. The Open Philanthropy Project is working on a report which explores the social cost of carbon in more detail.

10.3. Estimating climate damages

Quantifying the economic cost of climate change is very difficult. Modellers are faced with the challenge of predicting the damage that a range of climate impacts - on sea level rise, agriculture, heat effects on productivity, the effects on tropical diseases etc - will do over the course of several centuries.

There are two broad ways to calculate climate damages:⁴²²

⁴²² This framework is borrowed from Council of Economic Advisors, Office of Management and Budget, *Climate-Related Macroeconomic Risks and Opportunities*, 4 April 2022, https://www.whitehouse.gov/wp-content/uploads/2022/04/CEA_OMB_Climate_Macro_WP_2022.pdf. "Here we describe two strands of the literature most relevant to the task of incorporating physical damages into macroeconomic models: (1) "top-down" studies that attempt to estimate the effect of climate change on aggregate economic outcomes like GDP directly; and (2) "bottom-up" studies that assess the effects of climate change on a specific economic sector or category of damages, which can then be aggregated across categories to estimate total damages."

Top down studies that attempt to estimate the effect of climate change on aggregate economic outcomes, like GDP, directly.

Bottom up studies that assess the effects of climate change on a specific economic sector or category of damages, which can then be aggregated across categories to estimate total damages.

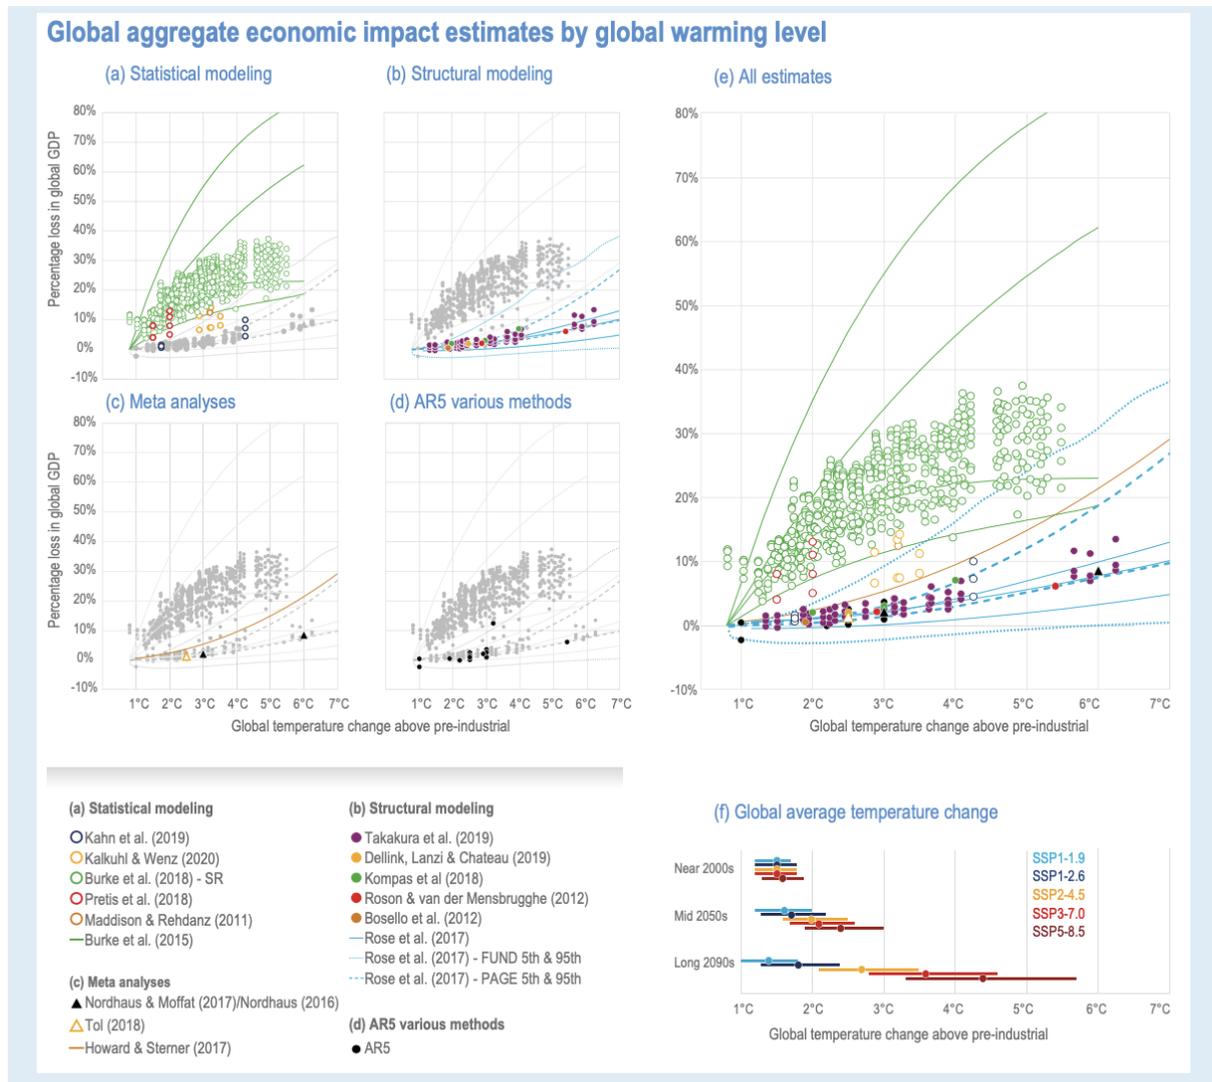
Hybrids of these two approaches are also possible.

One important difference between top-down and bottom-up approaches is that bottom-up approaches can in principle include non-market damages, i.e. damages that do not show up in GDP statistics. These include the intrinsic value of illness, death and the intrinsic value of ecosystems. A study that does not include non-market impacts might account for the *instrumental* but not intrinsic costs of illness or death, such as costs to the health system or in terms of lost productivity. But illness and death are important independently of their impact on GDP. It is better for people not to suffer or to die prematurely regardless of the effect this has on economic output.

These intrinsically bad non-market climate damages can be monetised and included in a damage function (though there is significant disagreement about how to do this and whether monetising these costs is appropriate rather than valuing them directly in units of welfare).⁴²³ The meaning of an estimate of 'climate damages as a fraction of GDP' will therefore be different depending on whether studies include non-market damages.

⁴²³ I briefly discuss some of these issues in this [blogpost](#).

The IPCC's Sixth Assessment Report provides the following useful diagram of damage estimates from different types of studies.



Source: IPCC, *Impacts*, Sixth Assessment Report, Ch 16, Figure Cross-Working Group Box ECONOMIC.1

Most studies find that for 4°C of warming, the monetised costs of climate change range from 5-10% of GDP. However, one set of top-down approaches which use statistical modelling to directly estimate impacts on GDP find much higher costs. The most prominent such study, Burke et al (2015), is shown in green in the chart above and finds that 4°C reduces global GDP by around 23%, and there is a 5% chance that it reduces global GDP by 60%. It is therefore very important to understand which of these studies are more plausible.

It is important to stress that these costs are relative to a counterfactual without climate change, not to today's level of income. On the low growth SSP3, global GDP per capita is not much higher than today, while on the high growth SSP5, it is nearly 1,000% higher. The other SSPs project that income per head will be several hundred percent higher. Thus, when a study says that, for instance, climate change of 4°C reduces GDP per capita by 10% on SSP2, it is saying that climate change causes average GDP per capita to be 270% higher, when it could have been 300% higher.

10.4. Bottom-up studies

Bottom-up studies assess the effects of climate change on a specific economic sector or category of damages, which can then be aggregated across categories to estimate total damages. Bottom-up studies will include many of the impacts I have discussed so far, including the effect on agriculture, sea level rise and the direct effects of rising heat on productivity.

10.4.1. Microeconomic evidence of how temperature affects productivity

Some microeconomic evidence suggests that rising temperatures might damage labour supply and labour productivity.

In Chapter 6 on heat stress, I discussed how climate change could affect labour supply. Several studies use industrial and military guidelines on safe working conditions to explore how rising heat stress might affect labour supply in different regions. The effect is summarised in the chart below from Buzan and Huber (2020), which assumes that people do not migrate and do not make use of extra air conditioning relative to today.

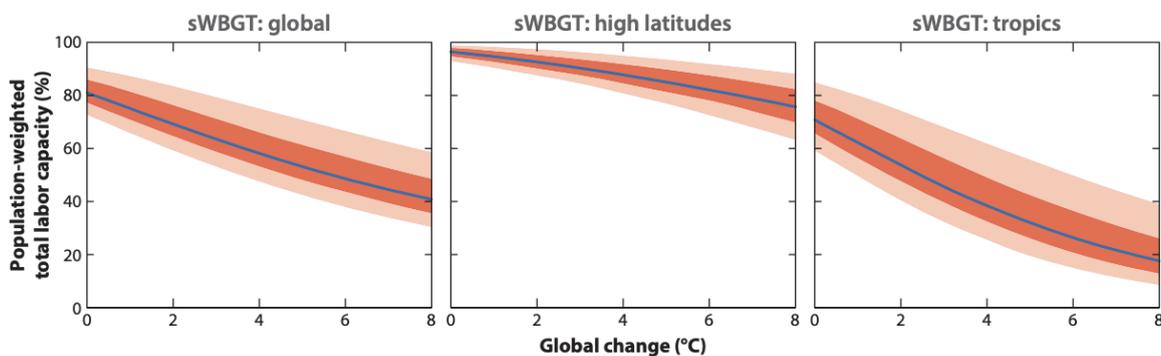


Figure 10

Population weighted total labor capacity. The CMIP5 ensemble is represented by the median (*blue line*), 50% (*red swath*), and 80% (*pink swath*) confidence intervals. The relative impacts on labor are shown at global (57°S to 57°N), high latitude (outside of 30°S to 30°N), and tropic (30°S to 30°N) regions. Abbreviations: CMIP, Coupled Model Intercomparison Project; sWBGT, simplified wet bulb globe temperature.

Source: Jonathan R. Buzan and Matthew Huber, 'Moist Heat Stress on a Hotter Earth', *Annual Review of Earth and Planetary Sciences* 48, no. 1 (2020): 623–55, <https://doi.org/10.1146/annurev-earth-053018-060100>.

Other studies also find that high temperatures reduce hours of labour supplied in industries with high exposure to climate.⁴²⁴

There is also microeconomic research suggesting that heat stress damages labour productivity per hour worked:

- Lab experiments suggest that there is a productivity loss in various cognitive tasks of about 2 percent per degree Celsius for temperatures over 25°C.⁴²⁵

⁴²⁴ Joshua Graff Zivin and Matthew Neidell, 'Temperature and the Allocation of Time: Implications for Climate Change', *Journal of Labor Economics* 32, no. 1 (2014): 1–26.

⁴²⁵ Dell, Jones, and Olken, 'What Do We Learn from the Weather?', 23.

- Observational and experimental studies also show a strong relationship between temperature and the productivity of factory, call centre, and office workers, as well as students. A meta-analysis of these studies concludes that increasing temperature from 23°C to 30°C reduces productivity by about 9 percent.⁴²⁶

The chart below from Heal and Park (2013) shows levels of performance at different temperatures. The optimal temperature for performance is around 22°C.

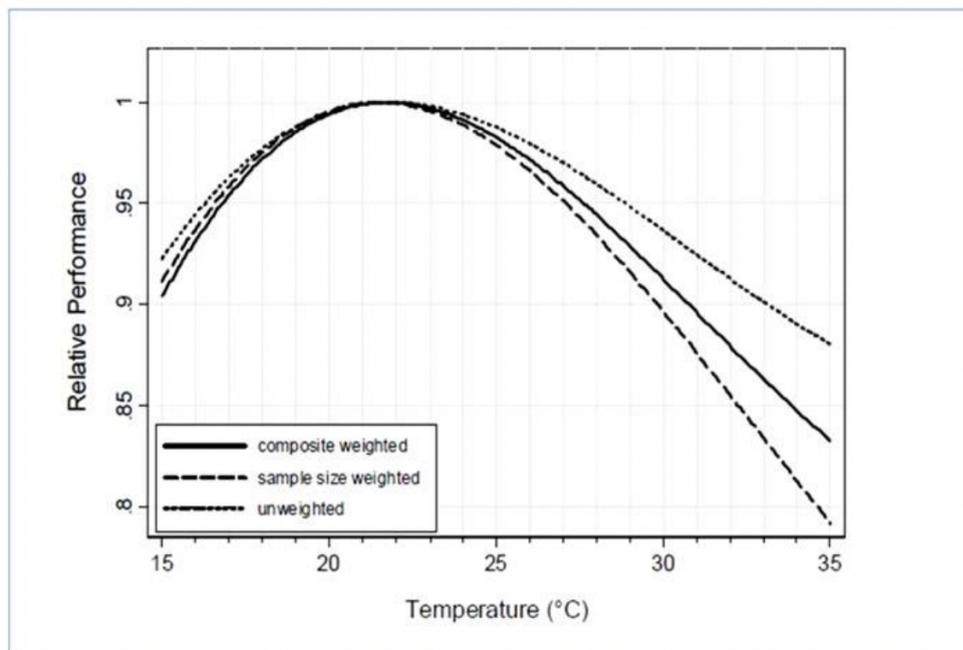


FIGURE 3.1. Task performance vs temperature. Maximum performance is normalized to 1 at 22 C. Source: Seppanen et al. [2006]

Source: Geoffrey Heal and Jisung Park, 'Feeling the Heat: Temperature, Physiology & the Wealth of Nations' (National Bureau of Economic Research, 2013).

Heal and Park (2013) find that air conditioning dampens the non-linear effects of temperature on economic performance, after imposing some controls.⁴²⁷

⁴²⁶ Dell, Jones, and Olken, 23.

⁴²⁷ "Consistent with the notion that higher levels of thermoregulatory capital dampen the impact of thermal stress on productivity, the subset of countries with above median air conditioning penetration feature a less concave relationship between temperature and income per capita. The temperature-income gradient implied by the coefficients on temperature and temperature squared in columns (26), (28), (30), and (32) – the subset countries with above-median air conditioning – is shallower than that implied by the coefficients in columns (25), (27), (29), and (31) – which represent the subset of countries with below-median air conditioning. Moreover, it seems that this difference is not being driven wholly by the correlation between air conditioning and other unobservables that are correlated with income. While countries with better access to thermoregulatory capital tend to be richer on average, there are also relatively hot and poor countries with high air conditioning penetration (for instance, Libya; see Table 7.6). It seems that the vulnerability to thermal stress as implied by access to thermoregulatory capital is not simply a function of "poorness" per se. This is an admittedly crude measure, but points us in the right direction for pressing policy-relevant research on climate adaptation." Heal and Park, 'Feeling the Heat'.

Rising heat stress would have the worst effects on the productivity of people working outdoors in industries like mining, construction and agriculture. This suggests that warming could be especially damaging for people in poor countries reliant on agriculture.

Today, agriculture accounts for [around 4% of global GDP](#), and around two thirds of the world economy is in services. There is more scope for people in the service sector to adapt to rising temperatures, both because they are indoors and because service-reliant economies tend to be richer. Insofar as workers are able to use air conditioning, they will be able to limit the damage to productivity from rising temperatures. The extent to which people will be able to make use of air conditioning depends on their socioeconomic prospects. The implications of different shared socioeconomic pathways for air conditioning penetration, according to one model, is shown below:

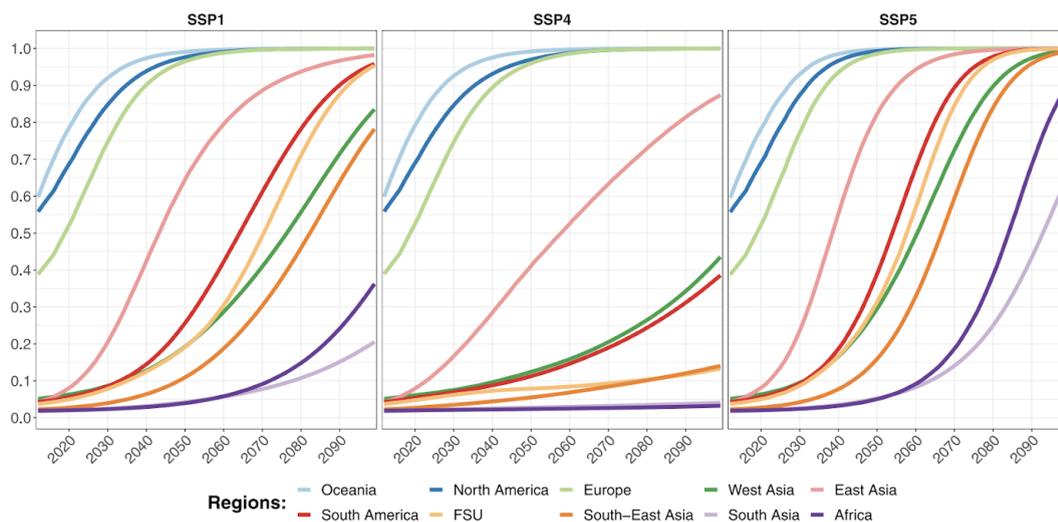


Fig. 4. Penetration rates of air conditioners by region and SSP.

Source: Orlov et al., ‘Economic Costs of Heat-Induced Reductions in Worker Productivity Due to Global Warming’.

I noted above that SSP4 is one of the more plausible SSPs. It suggests that the richest/highest growth regions - North America, Europe, Oceania and East Asia - will be able to reduce the effects of rising temperatures on productivity, while the low growth regions - South Asia and Africa - will have low take-up of air conditioning and so low ability to adapt to rising heat stress. Modellers disagree about how to model future air conditioning penetration, and some much more optimistic levels of penetration, on the SSPs.⁴²⁸

All in all, rising heat stress for workers is one plausible way that climate change could damage economic output.

⁴²⁸ Filippo Pavanetto et al., ‘Air-Conditioning and the Adaptation Cooling Deficit in Emerging Economies’, *Nature Communications* 12, no. 1 (9 November 2021): SI Fig. 5, <https://doi.org/10.1038/s41467-021-26592-2>.

10.4.2. Included sectors and impacts

The table below summarises the assumptions and findings of the bottom-up studies included in the above chart from the IPCC that summarises damage estimates in the literature.

Study	Sectors included	Non-market impacts?	Post-AR 5 literature ?	Tipping points?	Indirect effects ?	Cost of 4°C (% of GDP in 2100)
Roson and van der Mensbrugge (2012)	<ul style="list-style-type: none"> ● Sea level rise ● Variations in crop yields ● Water availability ● Human health ● Tourism ● Energy demand 	No ⁴²⁹	No	No	No	~5%
Bosello et al (2012)	<ul style="list-style-type: none"> ● Sea level rise ● Energy demand ● Agriculture ● Tourism ● Forests ● Floods ● Reduced work capacity due to heat stress 	No	No	No	No	~4%
Dellink et al (2017)	<ul style="list-style-type: none"> ● Sea level rise ● Crop yields ● Changes in fisheries catches ● Capital damages from hurricanes ● Temperature effects on labour productivity and human health ● Tourism ● Energy demand from heating and cooling 	No ⁴³⁰	No	No	No	~5%

⁴²⁹ “Changes in morbidity and mortality are interpreted as changes in labor productivity and demand for health care, and are used to shock exogenous parameters in a computable general equilibrium model, including 16 regions.” Roberto Roson and Dominique Van der Mensbrugge, ‘Climate Change and Economic Growth: Impacts and Interactions’, *International Journal of Sustainable Economy* 4, no. 3 (2012): 270–85.

⁴³⁰ “The health impacts of climate change have economic consequences that go beyond market costs. These costs, such as the costs of premature deaths, cannot be accounted for in the ENVLinkages model. However, they can be evaluated using WTP techniques and, for premature deaths, the Value of a Statistical Life.” Rob Dellink, Elisa Lanzi, and Jean Chateau, ‘The Sectoral and Regional Economic Consequences of Climate Change to 2060’, *Environmental and Resource Economics* 72, no. 2 (1 February 2019): 309–63, <https://doi.org/10.1007/s10640-017-0197-5>

Kompas et al (2018)	<ul style="list-style-type: none"> • Sea level rise • Losses in agricultural productivity • Temperature effects on labour productivity and human health • Energy demand • Tourism 	No ⁴³¹	No	No	No	~2%
FUND	<ul style="list-style-type: none"> • Sea level rise • Agriculture • Forestry • Water resources • Energy consumption • Human valuation of ecosystems • Human health • Extreme weather 	Yes ⁴³²	No	No	No	~5%
Takakura et al (2019)	<ul style="list-style-type: none"> • Fluvial flooding • Coastal inundation • Agriculture • Undernourishment • Heat-related excess mortality • Cooling/heating demand • Occupational-health costs • Hydroelectric generation capacity • Thermal power generation capacity 	Yes ⁴³³	Yes	No	No	~6%

As this shows, there is a fair amount of modeller discretion in bottom-up studies: different studies evaluate the effect on different sectors and use different literature to quantify these effects. As a rule, we should expect damage estimates to be biased low because they may have missed important sectoral impacts of climate change. None of the studies include the impact of tipping points or indirect effects. Only two of the models - the FUND integrated assessment model and Takakura et al (2019) - account for non-market impacts, like deaths and illness.

⁴³¹ “We aggregate and simulate labor productivity loss and human health damages via a negative labor productivity loss.” Tom Kompas, Van Ha Pham, and Tuong Nhu Che, ‘The Effects of Climate Change on GDP by Country and the Global Economic Gains from Complying with the Paris Climate Accord’, *Earth’s Future* 6, no. 8 (2018): 1153–73.

⁴³² FUND quantifies effects on health using the value of a statistical life, as discussed in the documentation [here](#).

⁴³³ “For undernourishment, heat-related excess mortality and fluvial flooding, the non-market values of lives lost represented by willingness-to-pay to avoid these risks were also incorporated.” Jun’ya Takakura et al., ‘Dependence of Economic Impacts of Climate Change on Anthropogenically Directed Pathways’, *Nature Climate Change* 9, no. 10 (October 2019): 737–41, <https://doi.org/10.1038/s41558-019-0578-6>.

Out of all of these models, Takakura et al (2019) seems like the most reliable because it incorporates a wide range of impacts which are commonly held to be important, it uses recent literature,⁴³⁴ and tries to account for non-market damages. The complete findings of Takakura et al (2019) are shown in the chart below

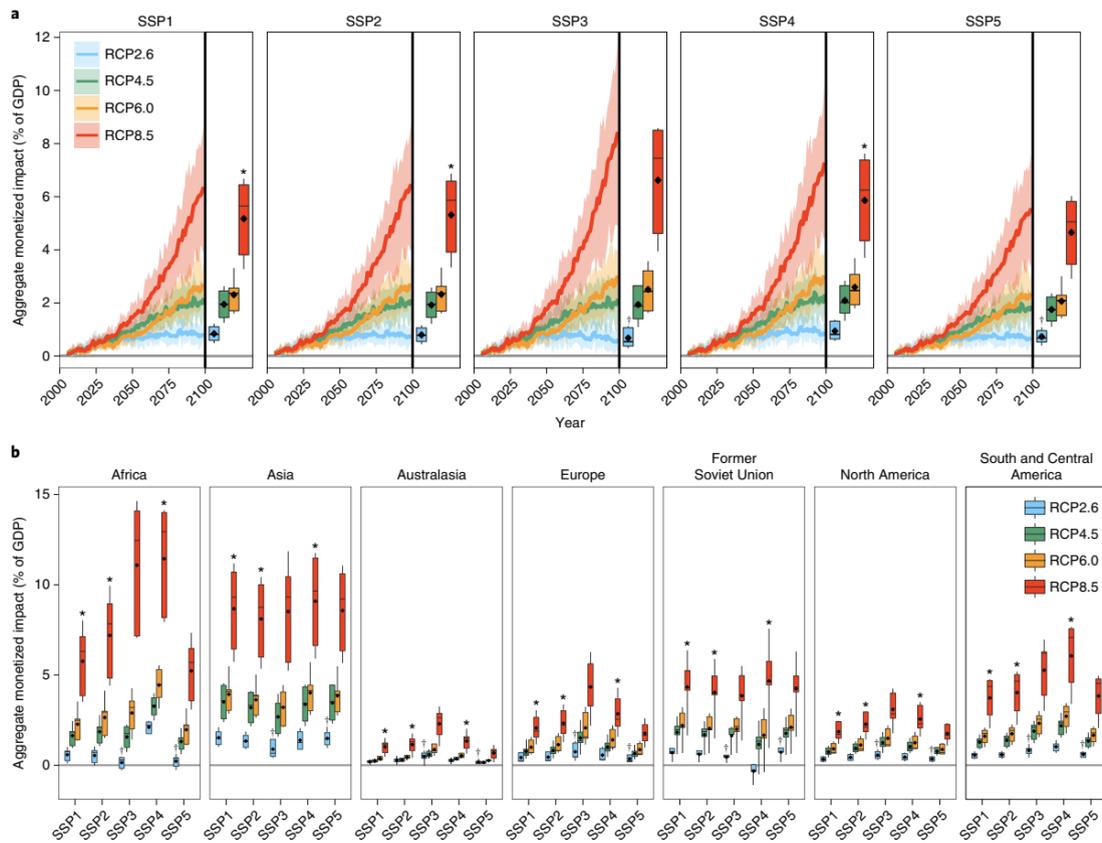


Fig. 1 | Projected economic impact. **a**, Time series of the aggregate monetized impacts and the mean impact for 2080–2099. **b**, Regional impacts. In **a**, the mean (line) and uncertainty range of the five GCMs (shaded area) are shown. For the boxplots in **a** and **b**, the mean impacts for 2080–2099 are shown and the range is the uncertainty range for the five GCMs. Centre line, median; dot, mean; box limits, upper and lower quartiles; whiskers, range. †Achieving the corresponding level of mitigation is less plausible under that SSP. *Reaching the corresponding level of emissions is less plausible under that SSP.

Source: Jun'ya Takakura et al., 'Dependence of Economic Impacts of Climate Change on Anthropogenically Directed Pathways', *Nature Climate Change* 9, no. 10 (October 2019): 737–41, <https://doi.org/10.1038/s41558-019-0578-6>.

On current policy, the chance of 4°C is around 5%, so the *expected* costs of the prospect of 4°C (i.e. weighting by the probability of the outcome) are around 0.4% of GDP.

However, Takakura et al (2019) excludes some important climate impacts, such as tipping points and indirect effects,⁴³⁵ and it also measures costs at aggregated regional scales which masks how uneven impacts are.

⁴³⁴ Their [supplementary information](#) shows that much of the literature cited is up to date.

⁴³⁵ "Although we covered most of the major sectors that will probably be affected by climate change, there also remain sectors such as conflict and crime, and the causal relationship between these sectors and climate change is controversial. The consequences of catastrophes were also not considered. Incorporation of these sectors is the next step in future work."

10.4.3. Tipping points

Most bottom-up studies exclude the impact of potential tipping points - low-probability high-impact events, such as the collapse of the West Antarctic Ice Sheet, the collapse of the Atlantic Meridional Overturning Circulation, or rapid warming from cloud feedbacks.⁴³⁶ Although these events are quite unlikely, their *expected* damage might be very high, so they probably account for a disproportionate share of the costs of climate change.

Dietz et al (2021) note that the treatment of tipping points in the literature is far from optimal

“A growing body of research has explored climate tipping points using economic models. We reviewed this literature and identified 52 papers that model the economic consequences of at least one climate tipping point (SI Appendix, Table S1). Many of these studies, however, represent climate tipping points in a highly stylized way. Examples include an instantaneous jump in the model's equilibrium climate sensitivity (11), an arbitrary reduction in global gross domestic product (GDP) (12), and a one-off permanent reduction in global utility (13).”⁴³⁷

Dietz et al (2021) estimate the economic cost of different tipping points by collecting together damages estimates for different tipping points from different integrated assessment models. They quantify the effect on GDP of the following tipping points:

- Permafrost carbon
- Methane hydrates
- Feedbacks from melting of Arctic sea ice
- Disintegration of the Greenland and West Antarctic Ice Sheets
- Slowdown of the Atlantic Meridional Overturning Circulation
- Variability of the Indian monsoon.

The chart below shows the effect that this has on the estimate of the social cost of carbon

⁴³⁶ “Another major problem with using IAMs to assess climate change policy is that the models ignore the possibility of a catastrophic climate outcome” Robert S. Pindyck, ‘Climate Change Policy: What Do the Models Tell Us?’, *Journal of Economic Literature* 51, no. 3 (September 2013): sec. 4, <https://doi.org/10.1257/jel.51.3.860>. See also Martin L. Weitzman, ‘GHG Targets as Insurance against Catastrophic Climate Damages’, *Journal of Public Economic Theory* 14, no. 2 (2012): 221–44.

⁴³⁷ Simon Dietz et al., ‘Economic Impacts of Tipping Points in the Climate System’, *Proceedings of the National Academy of Sciences* 118, no. 34 (2021).

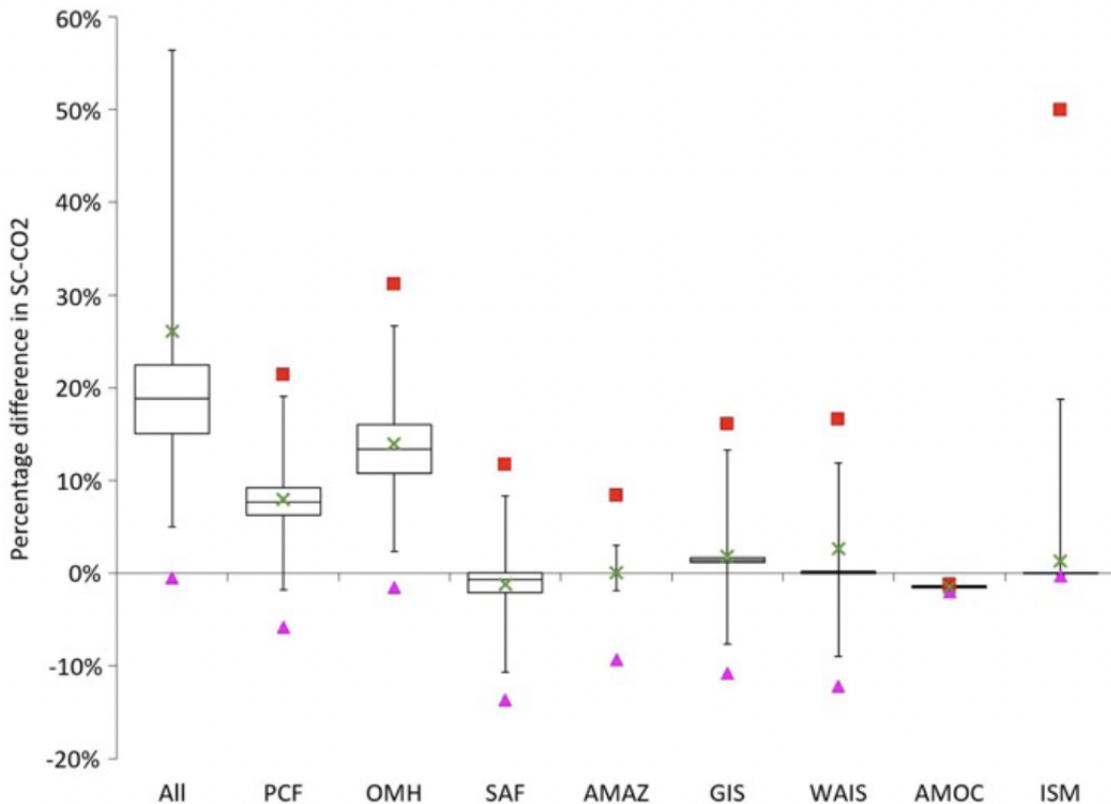


Fig. 1. The percentage change in the SCC due to tipping points collectively and individually. Boxes show medians and interquartile ranges, whiskers show 95% CIs, crosses mark the average changes (0.1% trimmed), triangles mark the 0.5 percentiles, and squares mark the 99.5 percentiles. The y axis is truncated. Specification comprises RCP4.5-SSP2 emissions and GDP/population growth, Hope and Schaefer PCF, Whiteman et al. beta OMH, and IPSL AMOC hosing. Monte Carlo sample size is 10,000.

Source: Simon Dietz et al., ‘Economic Impacts of Tipping Points in the Climate System’, *Proceedings of the National Academy of Sciences* 118, no. 34 (2021).

Tipping points increase the social cost of carbon by around 25% relative to no tipping points. They argue that accounting for non-market damages makes little difference to this estimate.⁴³⁸ Dietz et al (2021) acknowledge that their estimates are preliminary and likely too low.⁴³⁹

⁴³⁸ “SI Appendix, Fig. S20 and Table S13 report the effect of including a leading estimate of global nonmarket damages from climate change using the nonmarket damage module from the MERGE (Model for Evaluating Regional and Global Effects of GHG reductions policies) IAM (26). The resulting estimates of the SCC are more comprehensive but arguably more uncertain. The effect of all tipping points combined on the expected SCC increases marginally, to 26.9%” Simon Dietz et al., ‘Economic Impacts of Tipping Points in the Climate System’, *Proceedings of the National Academy of Sciences* 118, no. 34 (2021).

⁴³⁹ “Most of our numbers are probable underestimates, given that some tipping points, tipping point interactions, and impact channels have not been covered in the literature so far” Simon Dietz et al., ‘Economic Impacts of Tipping Points in the Climate System’, *Proceedings of the National Academy of Sciences* 118, no. 34 (2021).

Their estimates have been criticised for relying on arbitrary damage functions relating temperature to damages and for relying on unreliable literature.⁴⁴⁰ For instance, as the chart above shows, the model suggests that AMOC collapse would have net economic benefits, which doesn't seem plausible from the discussion in Chapter 8, which highlighted major transition costs from cooling and from precipitation changes, as well as weakening of the monsoons.

Given the lack of literature on the economic costs of tipping points, it may be useful to develop ballpark estimates of the costs of these tipping points informed by the scientific literature. In Chapters 7 and 8, I discussed the following tipping points:

1. Permafrost carbon
2. Methane clathrates
3. Amazon forest dieback
4. Boreal forest dieback
5. Collapse of the Atlantic Meridional Overturning Circulation
6. Changes in monsoons
7. Rapid cloud feedbacks
8. Collapse of the West Antarctic Ice Sheet

The key question is: what effect do these potential tipping points have on the estimate of the costs of climate change outlined in the previous subsection? I put most weight on Takakura et al (2019), which finds that for 4.4°C, climate damages are around 8% of GDP in 2100. The first four tipping points are carbon cycle feedbacks and so are not relevant to the costs of climate change at a given level of warming, though Amazon dieback would have large costs for Latin America.

AMOC collapse

AMOC collapse would cause cooling and drying around the North Atlantic. Britain would be particularly badly affected, and Ritchie found that AMOC collapse would reduce British GDP by around 0.3% due to irrigation spending. AMOC collapse would also affect global monsoons, with African and Asian monsoons weakening and southern hemisphere monsoons strengthening.⁴⁴¹ This could have large negative humanitarian consequences by causing drought and damaging agriculture. I am not sure how to quantify such costs.

We also need to weight these expected damages according to the probability of AMOC collapse. On business as usual warming of around 2.7°C, models suggest that the chance of collapse is well below 1%, while for 4°C (which has around 5% chance on business as usual), the chance is 1% to 5%, though models probably understate the risk. Thus, on current policy, the risk of AMOC collapse is probably upwards of (5%*5%=) 1 in 400. Since climate policy will probably strengthen in the future, the all-things-considered probability is

⁴⁴⁰ Steve Keen et al., 'Estimates of Economic and Environmental Damages from Tipping Points Cannot Be Reconciled with the Scientific Literature', Proceedings of the National Academy of Sciences 119, no. 21 (24 May 2022): e2117308119, <https://doi.org/10.1073/pnas.2117308119>.

⁴⁴¹ "A collapse of the Atlantic Meridional Overturning Circulation could weaken the African and Asian monsoons but strengthen the Southern Hemisphere monsoons (high confidence)." IPCC, Climate Change 2021: The Physical Science Basis, Sixth Assessment Report (UNFCCC, 2021), Box. TS.13.

lower. This suggests that the *expected* costs of AMOC collapse are at least one hundred times smaller than the actual cost of AMOC (because the badness of the outcome is weighted according to its probability).

Overall, if it were to occur, it seems plausible that AMOC collapse would have large costs, especially in Africa and Asia. However, given progress in emissions, the risk now seems low enough that it would not make a large difference to the bottom-up estimates of direct costs in Takakura et al (2019).

West African Monsoons

The IPCC suggests that greenhouse gases will probably cause there to be more rainfall in the central Sahel, but less in the Western Sahel. I am not aware of any models that try to quantify the effects this would have on social welfare. If the IPCC estimates are correct, there would be some benefits from increased rainfall in some regions, which would be offset by drying in other regions. The more concerning prospect seems to be weakening of the monsoons caused by AMOC collapse

Indian monsoon shift

The IPCC suggests that greenhouse gases will probably cause a strengthening of the Indian monsoon. I am not sure what the net effects of this would be on social welfare, but the more concerning prospect seems to be weakening of the monsoons caused by AMOC collapse

Collapse of the West Antarctic Ice Sheet

The collapse of the West Antarctic Ice Sheet would cause sea level rise of 5 metre over the course of around 100 years. Nicholls et al (2008) estimate that the costs of coastal protection for 5 metres of sea level rise quickly rise to \$30 billion per year, assuming perfect adaptation, which is around 0.04% of current global GDP.⁴⁴² Regional case studies of the Rhone, the Thames and the Netherlands suggest that there is greater potential for abandonment than if there were perfect adaptation.⁴⁴³ I am not aware of any estimates of the costs of coastal flooding assuming that adaptation measures are not taken.

Again, we need to weight the expected costs of these scenarios according to how probable they are. Models have mainly considered the risk of collapse of the West Antarctic Ice Sheet for 4.4°C, and there is around a 5% chance of that on current policy. The probability of collapse of the ice sheet conditional on 4.4°C is unclear, and the IPCC says that there is deep uncertainty about this possibility for more than 3°C of warming.⁴⁴⁴

⁴⁴² Nicholls, Tol, and Vafeidis, 'Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet', Fig. 5.

⁴⁴³ "While some observations of response to abrupt relative sea-level rise due to subsidence support the global model results, detailed case studies of the WAIS collapse in the Netherlands, Thames Estuary and the Rhone delta suggest a greater potential for abandonment than shown by the global model." Nicholls, Tol, and Vafeidis, 'Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet'.

⁴⁴⁴ IPCC, Climate Change 2021: The Physical Science Basis, Sixth Assessment Report (UNFCCC, 2021), Table 4.10.

Rapid cloud feedbacks

The effects of rapid cloud feedbacks outlined by Schneider et al (2019) would be much more damaging: there would be 8°C of warming over decades on top of around 5°C of warming. This would be extremely damaging, destroying agriculture in the tropics and subtropics and generally making those regions uninhabitable. The economic costs would be extremely large and I am not sure how to quantify them in a sensible way.

Given recent progress on climate policy, the probability of reaching the CO₂ levels which would trigger the cloud feedback - of 1,200ppm to 2,200ppm - now seem extremely unlikely. Burning all the fossil fuels would increase CO₂ concentrations to around 1,600ppm. In Chapter 1, I estimated that the chance of burning all of the fossil fuels was on the order of 1 in 500,000, and would take several centuries.

How do tipping points affect the economic costs of climate change?

Overall, the literature on the economic costs of tipping points is badly underdeveloped. Some of the tipping points seem like they would have very large costs if triggered, especially AMOC collapse, collapse of the West Antarctic Ice Sheet, and rapid cloud feedbacks. However, given progress on emissions, it looks increasingly unlikely that we will trigger these tipping points. If, as now seems likely, we limit warming to 2.7°C or below, the risk of the most damaging tipping points seems low enough that they would not make a dramatic change to the estimates from the bottom-up studies.

10.4.4. Indirect effects

According to Howard and Sterner (2017) most models do not include indirect effects such as conflict and crime.⁴⁴⁵ This biases damage estimates downward. I discuss these effects in Chapters 11-13.

10.4.5. Regional or global aggregation

In the main IAMs, known as DICE, FUND and PAGE, “much of the poverty associated with high levels of vulnerability is masked by regional averaging of economic variables”.

“For instance, the models used by the US Environmental Protection Agency to estimate the social cost of carbon [DICE (12), FUND (13), and PAGE (14); see ref. 15)] do not disaggregate below the level of continental regions. In particular, the entire population of each region is taken to consume the regional average. (And DICE does not disaggregate below the global level.)”⁴⁴⁶

Takakura et al (2019) aggregate the world into 17 regions.⁴⁴⁷

⁴⁴⁵ “Most of the underlying estimates systematically omit key climate impacts that could significantly increase climate damages, including socially-contingent climate impacts (migration, social and political conflict, and violence), ocean acidification, etc. (Howard 2014; Revesz et al. 2014).” Howard and Sterner, ‘Few and Not So Far Between’.

⁴⁴⁶ Francis Dennig et al., ‘Inequality, Climate Impacts on the Future Poor, and Carbon Prices’, *Proceedings of the National Academy of Sciences* 112, no. 52 (29 December 2015): 15827–32, <https://doi.org/10.1073/pnas.1513967112>.

⁴⁴⁷ “The impacts were originally calculated for a 0.5°×0.5° grid for each year during the twenty-first century and the calculated impacts were aggregated to 17 regions unless otherwise noted (above,

This has hugely important implications for the costs of climate change because there is significant inequality between countries within regions. Dennig et al (2015) modify a leading climate-economy model, Regional Integrated model of Climate and the Economy (RICE), to include what is known about economic inequality within regions and countries. Their findings are striking:

“When subregional differences are modeled in this way, several policy-relevant aspects of the model can change dramatically even when other assumptions and parameters from RICE are held constant. As we show below, even when RICE regional damage functions are used to establish the damage level of each region, the distribution of damage within regions can cause some members of future generations to be less affluent than their current counterparts. If the distribution of damage is less skewed to high incomes than the distribution of consumption, then weak or no climate policy will result in sufficiently large damages on the lower economic strata to eventually stop their welfare levels from improving, and instead cause them to decline. This paints a different picture from the standard narrative in leading cost–benefit IAMs, where regional average consumptions continue to grow even under business-as-usual (BAU).”⁴⁴⁸

This is important on utilitarian, egalitarian and prioritarian grounds. Standard climate-economy models suggest that everyone will be better-off despite climate change, but this is not realistic. As I discussed in the introduction to this report, this is one reason that discounting the future costs of climate change is inappropriate.

This is one area where presenting the overall costs of climate change as a fraction of GDP, rather than in terms of units of wellbeing lost, may cause confusion. A cost in terms of dollars is worth more to someone in Congo than someone in Norway. By presenting costs in terms of GDP, studies on damages mask this fact, and so understate the welfare costs of climate change. Once we correct for regional aggregation, even if we use Nordhaus’ damage model and his assumptions about discounting, it is optimal to limit warming to 2°C. In contrast, without regional disaggregation, it is optimal to limit warming 3.3°C. Regional disaggregation of impacts increases the optimal carbon price by a factor of 5 to 10.⁴⁴⁹ I’m not sure what effect this would have on the Takakura et al (2019) damage estimates.

10.4.6. Growth or levels?

Whether climate change affects the level of GDP or the growth rate is a crucial determinant of the size of climate damages. If temperature shocks merely have a level effect, then after damage in one bad year, the economy will bounce back in the next year. But if they have a growth effect, we would expect to see longer lasting and larger effects of temperature shocks.

results are further aggregated into seven regions).” Takakura et al., ‘Dependence of Economic Impacts of Climate Change on Anthropogenically Directed Pathways’.

⁴⁴⁸ Dennig et al., ‘Inequality, Climate Impacts on the Future Poor, and Carbon Prices’.

⁴⁴⁹ Dennig et al., ‘Inequality, Climate Impacts on the Future Poor, and Carbon Prices’, Fig. 1.

Stern and Stiglitz (2022) comment that in most integrated assessment models, economic growth is exogenous.⁴⁵⁰ Scholars have pointed to several mechanisms by which climate change might affect economic growth:⁴⁵¹

- Adaptation spending would divert spending away from R&D, which drives economic growth.⁴⁵²
- Morbidity, mortality and heat stress would damage technological progress and human capital.
- Damaged capital stocks

It is not clear how big an effect global warming would have via these mechanisms. However, my best guess is that the effects are small relative to the other determinants of growth. It is true that spending on adaptation would divert spending away from R&D, but this is also true of all government waste, and seems a weak and indirect lever on overall spending on R&D. There is evidence that heat stress damages productivity, which could damage innovation, but this effect can be removed with air conditioning, and the effects of plausible levels of warming seem small.

As I discuss below, some top-down studies try to quantify the effect on growth directly, though I am suspicious about some of the econometrics used. I would favour a bottom-up approach with a sophisticated causal model of the mechanisms by which climate change might impact growth, which is informed by the surrounding empirical literature on those mechanisms.

10.4.7. Overall judgement on bottom-up studies

Bottom-up approaches are, at first pass, a plausible way to calculate the costs of climate change. They can in principle use the scientific literature and sophisticated economic models to quantify and aggregate the most important climate impacts highlighted in the literature. The main drawbacks of extant bottom up studies are as follows:

1. Modeller discretion means that certain important sectors may be excluded or that scientific literature is not up-to-date.
2. No studies have well-developed treatment of tipping points.
3. Many studies quantify effects at high levels of regional aggregation without accounting for significant intra-regional inequality.
4. Most studies do not attempt to model potential effects on economic growth

⁴⁵⁰ “There is no damage to capital stocks in most IAMs, nor any reduction in the underlying growth rate, which is assumed to be exogenously determined.” Nicholas Stern and Joseph Stiglitz, ‘The Economics of Immense Risk, Urgent Action and Radical Change: Towards New Approaches to the Economics of Climate Change’, *Journal of Economic Methodology*, 2022, 1–36.

⁴⁵¹ Jarmo S. Kikstra et al., ‘The Social Cost of Carbon Dioxide under Climate-Economy Feedbacks and Temperature Variability’, *Environmental Research Letters* 16, no. 9 (September 2021): 094037, <https://doi.org/10.1088/1748-9326/ac1d0b>.

⁴⁵² “Moreover, the high levels of expenditure necessary to adapt to climate change, especially in the more adverse scenarios, implies that resources will be diverted away from innovation and investment and the social cost of carbon in those states of nature will be high.” Stern and Stiglitz, ‘The Economics of Immense Risk, Urgent Action and Radical Change’.

Most bottom-up models find that the monetised costs of warming of 4°C are equivalent to a decline in GDP of around 5% in 2100. The most plausible model I have looked into, Takakura et (2019), finds that the monetised costs of 4.4°C are equivalent to around a 6% reduction in GDP by 2100.

Due to the issues outlined above, this is very likely an underestimate of the true costs of climate change. Known tipping points seem important for higher levels of warming, but the probability of 3.5°C or more is now low enough that they seem unlikely to make a dramatic difference to the expected costs of climate change. However, there may be unknown tipping points which are more important. The literature highlights that the effects would be disproportionately borne by the worst-off people, who have contributed the least to climate change.

One reason that estimates in the literature may be overestimated is that they may not accurately capture the scope for adaptation of future richer societies, which is hard to predict, but may soften the blow of climate change. Still, I think that impacts are, on the whole, probably overestimated.

It is difficult to see how plausible changes in the estimates of direct damages found in bottom-up studies could come close to threatening a global catastrophe, or the collapse of industrial civilisation. Even once we correct for the problems outlined above, it still seems like average global living standards will improve. It is very clear from the bottom-up studies that strong mitigation is justified, but it is also difficult to see how plausible levels of climate change could cause civilisational catastrophe.

I discuss the import of indirect effects in Chapters 11 to 13.

10.5. Top-down studies

Top down studies attempt to estimate the effect of climate change on aggregate economic outcomes, like GDP, directly. Much of the recent literature that directly estimates the effects of climate change on GDP uses panel data and short-term variations in weather to estimate the relationship between climate variables and GDP based on past experience.⁴⁵³

Heuristically, on this methodology, an economy observed during a cool year is the ‘control’ for that same society observed during a warmer ‘treatment’ year. Because variation in weather is exogenous, any effects on output are likely to be causal. These studies can also explore the effects of weather variation on shorter subannual timescales.

Some of these studies have produced markedly higher damage estimates than bottom-up studies. Before I discuss these studies it is worth exploring the correlation between climate variables and economic outcomes across space.

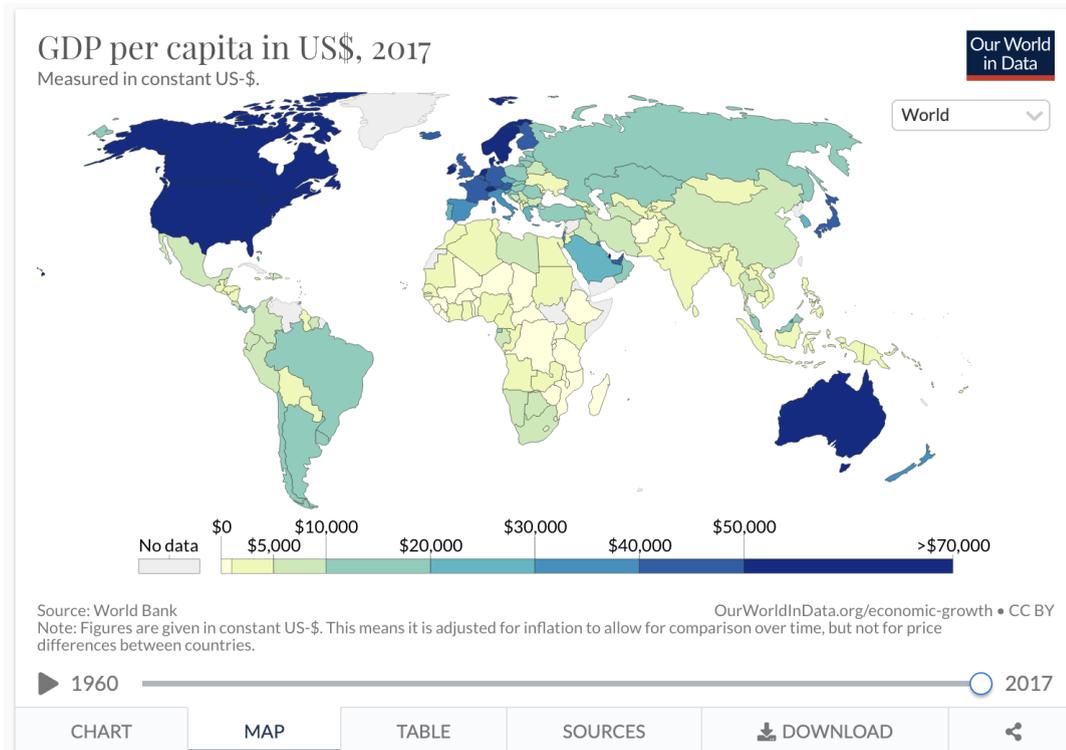
⁴⁵³ Council of Economic Advisors, Office of Management and Budget, *Climate-Related Macroeconomic Risks and Opportunities*, 4 April 2022, p. 9, https://www.whitehouse.gov/wp-content/uploads/2022/04/CEA_OMB_Climate_Macro_WP_2022.pdf

10.5.1. Cross-sectional correlations

Cross-country correlations

The cross-country correlation between temperature and per capita GDP

There is a strong negative relationship between hot climates and per capita GDP. One can see this just from eyeballing a map of income per head



Latitude is an imperfect proxy of temperature. Gallup, Sachs and Mellinger (1999) show that countries located in the tropics (i.e. between the Tropic of Cancer and the Tropic of Capricorn) were 50% poorer per-capita in 1950 and grew 0.9 percentage points slower per year between 1965 and 1990.⁴⁵⁴

This chart plots ln GDP per capita against population-weighted average temperature:

⁴⁵⁴ John Luke Gallup, Jeffrey D. Sachs, and Andrew D. Mellinger, 'Geography and Economic Development', *International Regional Science Review* 22, no. 2 (1 August 1999): 179–232, <https://doi.org/10.1177/016001799761012334>.

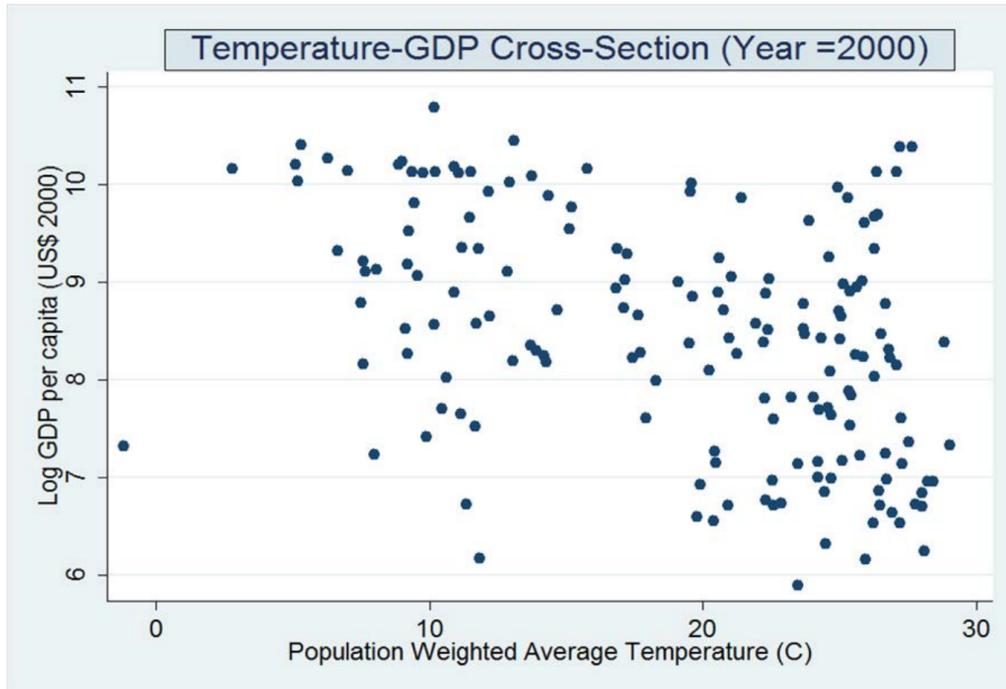


FIGURE 2.1. Countries by log income per capita and population-weighted average temperature

Source: Geoffrey Heal and Jisung Park, 'Feeling the Heat: Temperature, Physiology & the Wealth of Nations' (National Bureau of Economic Research, 2013).

For reference, $\ln(\$60,000) = 11$; $\ln(\$8,000) = 9$; $\ln(\$400) = 6$.

This masks the strength of the correlation because it puts income on a log scale. Still, it is clear that at a given temperature, there is a fair bit of variation.

As we saw in the section on heat stress, temperature *and humidity* contribute to heat stress, which in turn affects productivity. There is also evidence that higher heat stress is correlated with lower GDP per capita:

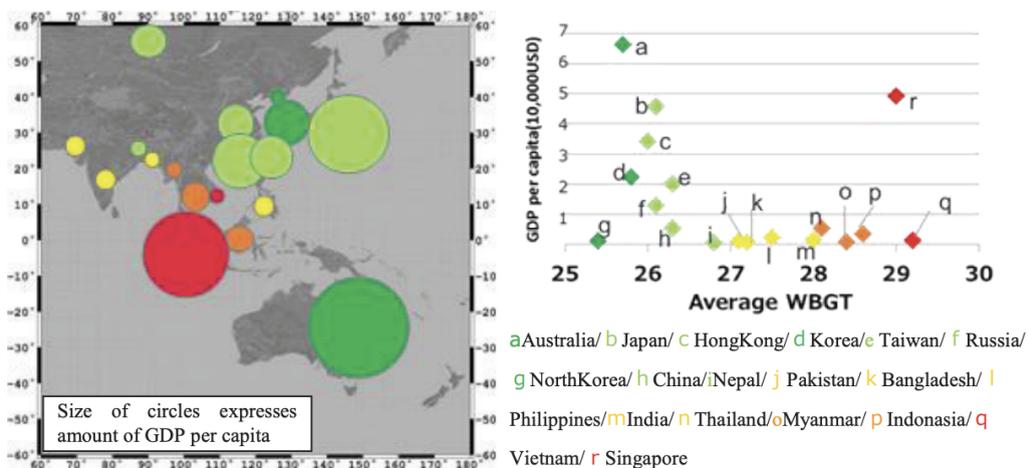
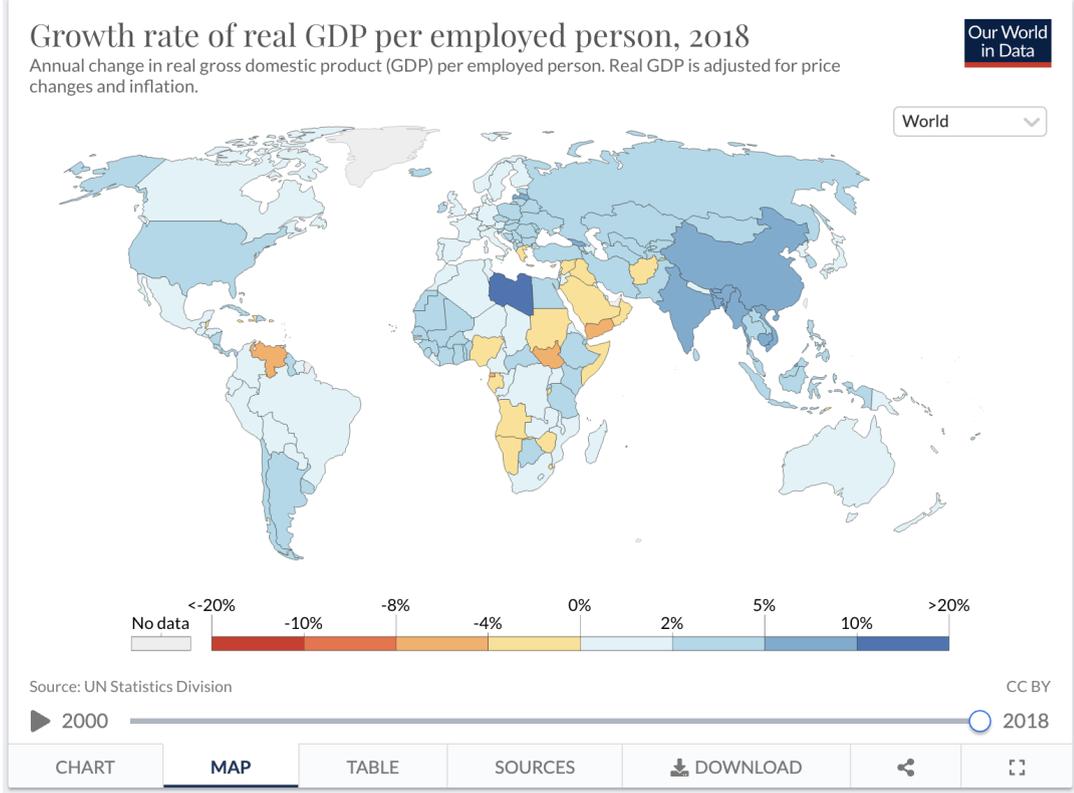


Figure 5. Relation between WBGT and GDP per capita.

Source: Noriko Okamura et al., 'Evaluating Thermal Comfort in City Life and Its Relation to Socio-Economic Activities', *Asian Journal of Geoinformatics* 14, no. 2 (2014).

The cross-country correlation between temperature and growth rates

The 'growth *rate*-tropics' correlation is less clear.



Here, the correlation doesn't seem especially strong. Most notably, in the last few decades, countries in East and South Asia have recently had much higher rates of growth than the West despite being much warmer on average. The same is true for several countries in Sub-Saharan Africa. The table below shows growth rates in selected large countries from 2000 up until before COVID-19.

Country	Average growth rate in GDP per capita
Vietnam	>5%
India	>4%
Bangladesh	>4%
Indonesia	~4%
Rwanda	>4%
Ethiopia	>5%

Source: [World Bank](#)

Methodological issues with cross-country regressions

The correlations I have outlined above do not necessarily suggest that there is a causal relationship between climate and economic indicators. Identifying a causal relationship from cross-sectional evidence is difficult for several reasons.

Firstly, there is *omitted variable bias*. There are multitudinous other factors that might be correlated with climate that could influence economic performance, and it will be difficult to identify and control for all of them properly. For example, in *WEIRDest People in the World*, Joseph Henrich argues that historical rates of cousin marriage affect growth rates,⁴⁵⁵ but I would wager that few cross-sectional studies have thought to control for that.

Furthermore, controlling for potential confounding variables, such as institutions or population, risks overcontrolling:

“For example, consider the fact that poorer countries tend to be both hot and have low-quality institutions. If hot climates were to cause low quality institutions, which in turn cause low income, then controlling for institutions... can have the effect of partially eliminating the explanatory power of climate, even if climate is the underlying fundamental cause.”⁴⁵⁶

Secondly, regressions of climate on economic performance might identify long-run historical effects that are not applicable to future climate change. Suppose that, as some scholars believe, temperature and income are correlated in the cross section today largely because climate affected the path of agricultural development, technological exchange, and/or subsequent colonialism. The cross-sectional relationship, which represents a very long-run equilibrium, may incorporate processes that are too slow to accurately inform the time-scale of interest, or it may include historical processes (such as colonialism) that will not repeat themselves in modern times.⁴⁵⁷

To illustrate, according to Acemoglu et al (2001), the key factor that explains the success of economies is that patterns of colonialism and institutions were influenced by mortality rates of settlers between the 16th and 19th century.⁴⁵⁸ These mortality rates would have been affected by local climate at the time because climate variables influence disease burden. This would explain why a hot climate is associated with poor economic performance today. However, this does not indicate that future warming will damage economic output in the future since colonialism won't repeat itself.⁴⁵⁹

⁴⁵⁵ Joseph Henrich, *WEIRDest People in the World* (Picador Paper, 2021).

⁴⁵⁶ Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken, 'What Do We Learn from the Weather? The New Climate-Economy Literature', *Journal of Economic Literature* 52, no. 3 (2014): 6.

⁴⁵⁷ Dell, Jones, and Olken, 6.

⁴⁵⁸ Daron Acemoglu, Simon Johnson, and James A. Robinson, 'The Colonial Origins of Comparative Development: An Empirical Investigation', *American Economic Review* 91, no. 5 (December 2001): 1369–1401, <https://doi.org/10.1257/aer.91.5.1369>.

⁴⁵⁹ Acemoglu (2001) is a flawed study. Firstly, the data on settler mortality have been called into question by Albouy (2012). Secondly, to accept the instrumentation in Acemoglu et al (2001), one has to assume that disease burden in the countries of interest only had an effect on growth via influencing settler mortality and the institutions they set up, and not through any other causal mechanism. But perhaps having lots of malaria in your country is bad for growth for other reasons. David Y. Albouy, 'The Colonial Origins of Comparative Development: An Empirical Investigation: Comment', *American Economic Review* 102, no. 6 (May 2012): 3059–76, <https://doi.org/10.1257/aer.102.6.3059>.

In light of these methodological difficulties, cross country-regressions are now out of fashion in economics, and the discipline has moved towards studies that can more reliably identify causal relationships.⁴⁶⁰

In my view, throwing out all evidence from cross-country correlations is too drastic; although they should be treated with caution, I think in some cases cross-sectional correlations provide useful information. The reason for this is that in some cases, economic theory and evidence from other domains may give us reason to believe that a particular relationship is in fact causal. Just because a relationship *might* be confounded does not mean that it is confounded.

For example, there is a very strong cross-sectional correlation between smoking and lung cancer: smokers are much more likely to get lung cancer than non-smokers. It is true that this relationship might be confounded by various demographic factors that independently cause cancer: people of lower socioeconomic status are more likely to smoke, so maybe their higher rates of lung cancer are driven by other aspects of their lifestyle. But even before we do a randomised control trial on the effects of smoking, or gather longitudinal evidence on the timing of increases in lung cancer rates and smoking, the very strong correlation between smoking and lung cancer should update us towards the view that one causes the other. This is especially true because we have an independently plausible theoretical explanation of why smoking causes lung cancer.

In sum, cross-sectional shouldn't be thrown out entirely but should still be treated with significant caution.

One alternative is to combine cross-sectional and panel data approaches.⁴⁶¹

Cross-sectional relationship within countries

One of the most important potential confounders of the cross-country relationship between temperature and income is country-level fixed effects, such as policies, state capacity and institutional quality. This is because the policies and institutions of countries are widely thought to be a major determinant of the economic prospects of different countries. This can be seen by looking again at recent growth rates in warmer countries. Recent growth

⁴⁶⁰ Durlauf, Steven N. 2000. "Econometric Analysis and the Study of Economic Growth: A Skeptical Perspective." In *Macroeconomics and the Real World: Volume 1: Econometric Techniques and Macroeconomics*, edited by Roger E. Backhouse and Andrea Salanti, 249–62. Oxford and New York: Oxford University Press.

⁴⁶¹ "A second approach along these lines represents a marriage of the panel data estimation approach using short-run weather fluctuations and the Ricardian approach. The concept here is that if one observes a large number of units (like counties, households, or firms) over a significant number of periods covering a spatial area with large heterogeneity in climate, one can estimate separate response functions for subgroups of the individual units using observed short-run weather fluctuations (for example, use within-household variation to identify a short-run response function by zip code). By controlling for unit- and time-fixed-effects, it is possible to obtain plausibly causal estimates of local short-run dose response functions. One can then either in a second step regress the slopes of the dose response on climate (for example, long run average summer temperature) across subgroups, or, through an interaction term in a single regression, estimate how the slope of the dose response function varies across areas with different climates, incomes, and other observables that vary across space." Maximilian Auffhammer, 'Quantifying Economic Damages from Climate Change', *Journal of Economic Perspectives* 32, no. 4 (November 2018): 33–52, <https://doi.org/10.1257/jep.32.4.33>.

accelerations in South and East Asia were caused by changes in policy, not by changes in climate.

One way to control for country-level fixed effects is to examine the cross-sectional relationship between climate and economic output *within* countries. For example, there is substantial variation in annual average temperature across US states, with the hottest state nearly 17°C warmer than the coldest state.

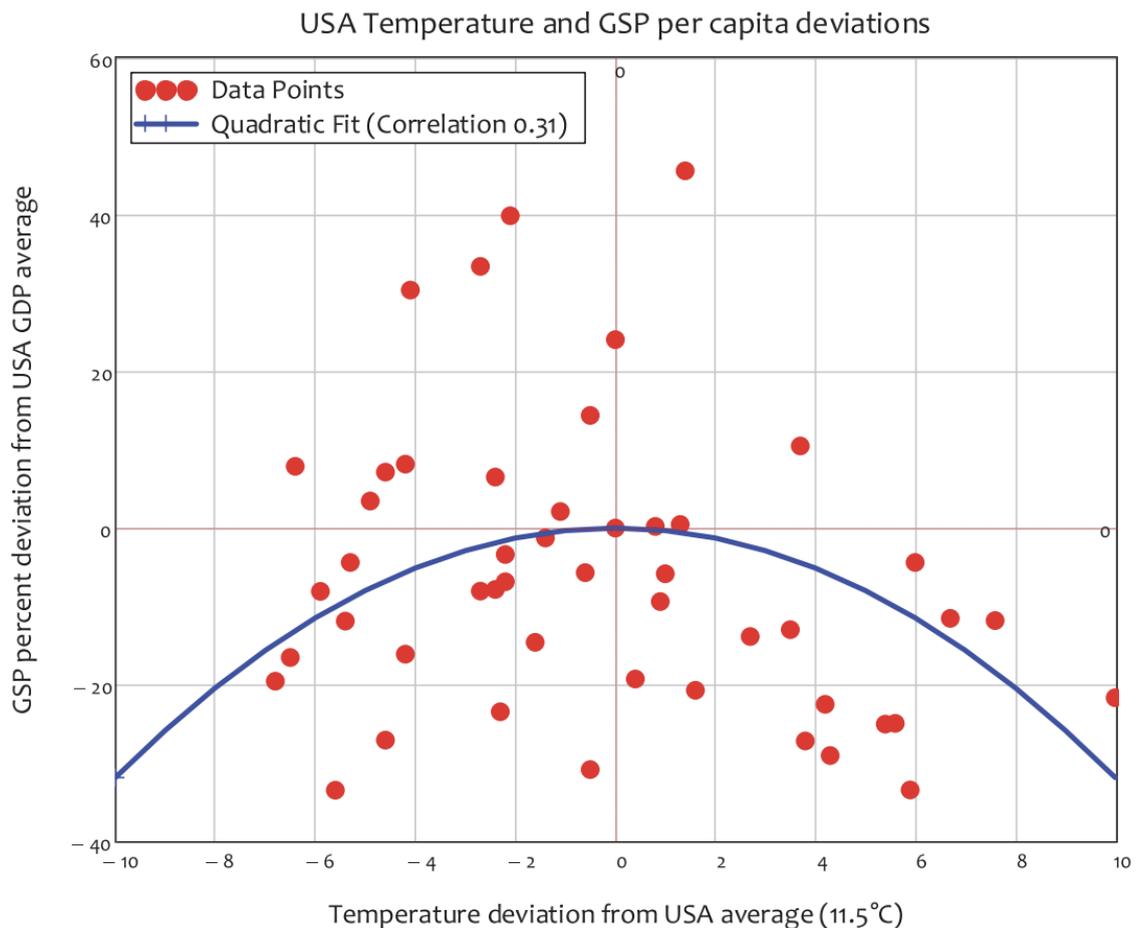


Figure 2. Correlation of temperature and USA Gross State Product per capita.

Source: Steve Keen, 'The Appallingly Bad Neoclassical Economics of Climate Change', *Globalizations* 0, no. 0 (1 September 2020): 1–29, <https://doi.org/10.1080/14747731.2020.1807856>.

In spite of this variation, in global and historical context, all of these states are very rich.

The findings of some major studies taking a subnational cross-sectional approach are summarised in the table below:

Study	Findings
Nordhaus (2006) ⁴⁶²	Controlling for country fixed effects, this study finds that 20% of the

⁴⁶² William D. Nordhaus, 'Geography and Macroeconomics: New Data and New Findings', *Proceedings of the National Academy of Sciences* 103, no. 10 (7 March 2006): 3510–17, <https://doi.org/10.1073/pnas.0509842103>.

	income differences between Africa and the world's rich industrial regions can be explained by geographic variables, which include temperature and precipitation as well as elevation, soil quality, and distance from the coast
Dell et al (2009) ⁴⁶³	A 1°C rise in temperature is related to a 1.2-1.9% decline in municipal incomes for 7,684 municipalities in 12 countries in the Americas. The within-country cross-sectional correlation is substantially weaker than any cross-country correlation.
Greßer et al (2021) ⁴⁶⁴	No negative relationship between subnational temperature and four different measures of economic development (per capita GDP, growth of per capita GDP, nightlights and gross cell production). There is no evidence that temperature is non-linearly related to income (with hotter regions being potentially particularly prone to adverse effects of temperature on income). There is no robust evidence that the effect of temperature is especially pronounced in poorer regions.

A key drawback of within-country cross sectional correlations is that they cannot shed light on how climate change might affect the choice of policies and institutions between different countries, which is an important limitation.

Nonetheless, they do shed light on some forms of climate impact. Most importantly, the effects of temperature on labour productivity should show up in within-country cross-sectional regressions. The lack of effect found in Greßer et al (2020) should update us towards the effect being small.

Methodological issues with within-country cross-sectional regressions

Within-country cross-sectional studies are also subject to omitted variable bias and to over-controlling. It might be that past climate is correlated with some feature which makes one subnational region thrive and another not. For instance, climate is correlated with rates of cousin marriage in different regions in Italy, which, according to Joseph Henrich and others, contributes to southern Italy's corruption, mafia problem, and poor economic performance. But this is a causal product of the willingness of the Church centuries ago to prohibit cousin marriage in northern Italy, which is not relevant to the future impact of climate change.

Although the within-country correlation might be confounded for specific countries, it is hard to see why it would be systematically confounded across all countries.

Other problems with cross-sectional evidence

Some aspects of future climate change are not well-captured by current variation in climate variables across space. This includes things like sea level rise, ocean acidification and CO₂ fertilisation.

⁴⁶³ Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken, 'Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates', *American Economic Review* 99, no. 2 (2009): 198–204.

⁴⁶⁴ Christina Greßer, Daniel Meierrieks, and David Stadelmann, 'The Link between Regional Temperature and Regional Incomes: Econometric Evidence with Sub-National Data', *Economic Policy* 36, no. 107 (2021): 523–50.

Since they measure the effect of climate variables on GDP directly, they also exclude non-market impacts, such as illness and death.

Overall judgement on the cross-sectional evidence

My conclusions about the cross-sectional evidence are as follows.

Firstly, I don't put much weight on the current cross-country relationship between temperature and GDP levels, for several reasons. The sample of countries on which to carry out a regression is not particularly large, which limits statistical power. There are also numerous potential confounds, which it is very difficult to properly control for.

Secondly, in the longer-term, the climate-growth *rates* relationship is more important than the climate-levels relationship. The correlation between growth rates and temperature does not seem particularly strong. If the various countries in the tropics and subtropics experiencing high economic growth continue to do so, then they will catch up with high-income countries within a few decades. It is a mistake to categorise the current relative poverty of Bangladesh as an instance of climatic determinism. Given their current growth rates, they are likely to be high-income countries soon, and that is due to policy change, rather than climate.

The current global variation in economic growth rates suggests that they are mainly determined by economic policy. Indeed, economic growth has mainly been a phenomenon of what Lant Pritchett calls the post-1950 '[development era](#)', including

- The end of colonisation with the liberation of India, Pakistan and Indonesia
- The founding of the Bretton Woods institutions - the IMF and the World Bank
- Truman's Four Point plan to provide technical assistance to developing countries
- Overall a concerted effort by economists and sovereign states to increase development

Growth has only been an independent area of study since this point. This era has brought more progress than all prior human history *combined*. Consequently, many countries that are currently hot will be much richer in the future.

Finally, the within-country cross-sectional regressions shed light on the effects of temperature on labour productivity.

However, they also miss several avenues of climate impact:

1. How climate affects the choice between different institutions and policies.
2. The effects of precipitation change
3. The transition costs of climate change
4. The costs of tipping points
5. The costs of longer-term impacts such as
 - a. Sea level rise
 - b. Ocean acidification
 - c. CO₂ fertilisation
6. Non-market impacts like illness and death

Overall, this suggests that we should put much more weight on bottom-up studies than on cross-sectional studies.

10.5.2. Weather and economic performance

Seminal studies by Dell et al (2012) and Burke et al (2015a) have used panel data to estimate the effect of inter-annual changes in weather on economic performance within countries, and used these to estimate the potential economic effects of climate change.⁴⁶⁵ After a large number of robustness checks, both papers conclude that there is no evidence that rainfall has a significant and consistent impact on GDP growth, but higher temperatures have adverse consequences. The damage estimate of Burke et al (2015a) is far higher than that of Dell et al (2012). As I discuss below, this is mainly due to the controls used in the different studies.

Heuristically, on this methodology, an economy observed during a cool year is the ‘control’ for that same society observed during a warmer ‘treatment’ year. Because interannual weather change is exogenous, any difference in economic output is plausibly due to differences in weather rather than country-specific or time trend factors.

This methodological approach avoids the omitted variables and over-controlling problems associated with cross-sectional approaches. A further advantage of the top-down approach is that it captures feedback between different economic sectors and damage categories, which are not well-captured when economic sectors and damage categories are estimated through a bottom-up approach and then simply aggregated.⁴⁶⁶

Methodological drawbacks of weather studies

This approach also has some drawbacks.⁴⁶⁷

Weather is different to climate change

Inter-annual weather change is importantly different to long-term climate change. The chart below shows the year-to-year change in global average surface temperature between 1950 and 2010.

⁴⁶⁵ Marshall Burke, Solomon M. Hsiang, and Edward Miguel, ‘Global Non-Linear Effect of Temperature on Economic Production’, *Nature* 527, no. 7577 (November 2015): 235–39, <https://doi.org/10.1038/nature15725>; Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken, ‘Temperature Shocks and Economic Growth: Evidence from the Last Half Century’, *American Economic Journal: Macroeconomics* 4, no. 3 (July 2012): 66–95, <https://doi.org/10.1257/mac.4.3.66>.

⁴⁶⁶ Council of Economic Advisors, Office of Management and Budget, *Climate-Related Macroeconomic Risks and Opportunities*, 4 April 2022, p. 10 https://www.whitehouse.gov/wp-content/uploads/2022/04/CEA_OMB_Climate_Macro_WP_2022.pdf.

⁴⁶⁷ Dell, Jones, and Olken, ‘What Do We Learn from the Weather?’, 37ff; Council of Economic Advisors, *Climate-Related Macroeconomic Risks and Opportunities*, 2022, p. 10-11.

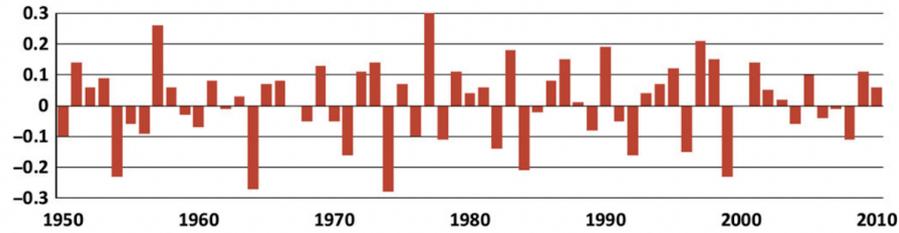
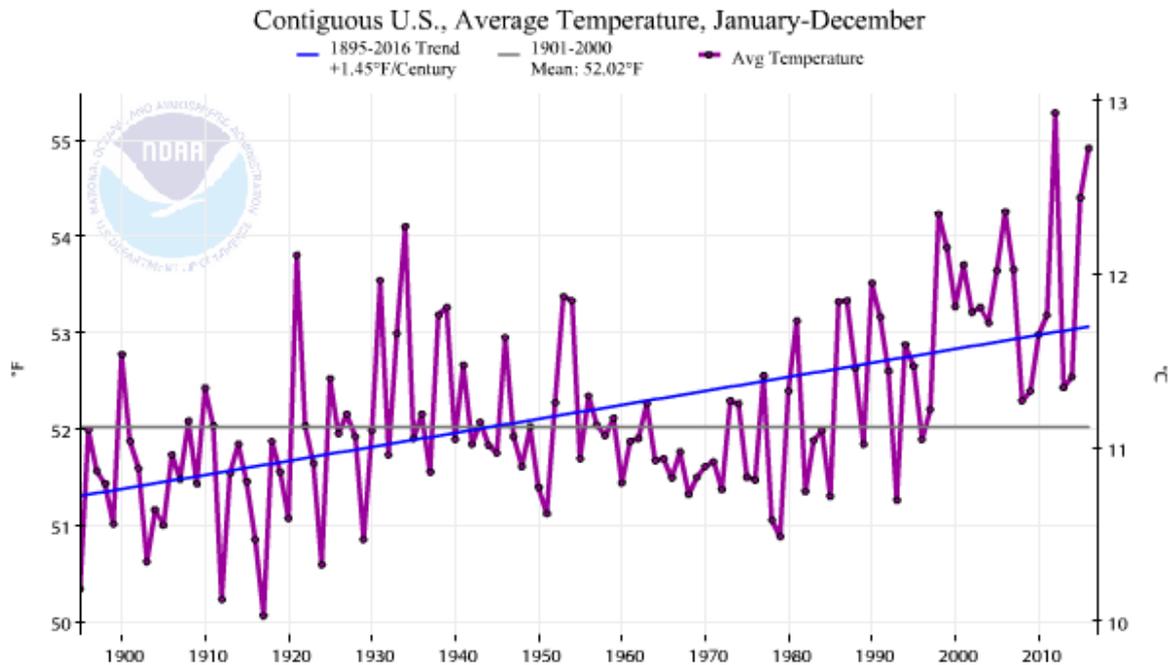


FIGURE 2 Year-to-year change in global mean surface temperature, 1950–2010 (vertical scale is °C; from <http://www.bom.gov.au/climate/change/>)

Source: John CV Pezzey, 'Why the Social Cost of Carbon Will Always Be Disputed', *Wiley Interdisciplinary Reviews: Climate Change* 10, no. 1 (2019): e558.

The maximum year-to-year change in global mean surface temperatures during 1950–2010 was only 0.3°C, which is much smaller than projected climate change.

The chart below is another way to illustrate the difference between climate change and interannual weather variation:



Between 1997 and 1998, the annual average temperature jumped 1°C in the US. But climate change is occurring at a rate of about 0.02°C per year. The weather studies measure the effect of year-to-year shifts in the purple line, whereas climate change is very slowly shifting the distribution of the purple line up over time, which is quite a different effect.

This could introduce bias in either direction. On the one hand, this misses potential tipping points, as well as effects that do not see sharp shifts on inter-annual timescales, such as ocean acidification, CO₂ fertilisation, and sea level rise.

On the other hand, we have much more time to adapt to slow moving climate change than we do to changes in average annual temperatures. There would also be more time for

general equilibrium effects, such as the movement of labour and capital, to reduce the damage of climate change.⁴⁶⁸

Noisy proxies

Proxy measures for temperature and precipitation will be quite noisy measures of the true weather signal. As a covariate becomes noisier, this increases the risk of finding a false positive in a regression.⁴⁶⁹ As Auffhammer (2018) says

“Another critique of the panel data approach is that if weather is measured with error, then as more fixed effects are included in the regression, concerns over measurement error loom larger (Fisher, Hanemann, Roberts, and Schlenker 2012). In the vast majority of locations, weather is measured with error, and the bigger the distance between weather stations, the bigger measurement error concerns become. The United States and Europe have tens of thousands of weather stations, but many locations in sub-Saharan Africa do not have a weather station within hundreds of miles. If the measurement error is classical, this is likely to attenuate the response towards zero.”⁴⁷⁰

This problem also applies to within-country cross-sectional approaches,⁴⁷¹ though to a more limited extent because studies can reduce noise in regional temperature data by averaging across long periods of time. For instance, in their within-country cross sectional regressions, Greßer et al (2021) measure average subnational temperature between 1950 and 2000.⁴⁷²

Non-market impacts

Since top-down panel studies try to measure the effect of climate variables on GDP directly, they exclude non-market impacts such as illness and death.

⁴⁶⁸ “If both labor and capital are mobile, then this type of macroeconomic readjustment could reduce the long-run impacts of climate change relative to a short-run panel estimate (although any such tempering of the impacts would depend on moving costs, the extent to which the marginal product of capital is location specific, and potentially a host of other factors).” Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken, ‘What Do We Learn from the Weather? The New Climate-Economy Literature’, *Journal of Economic Literature* 52, no. 3 (2014): 740–98.

⁴⁶⁹ “By contrast, multiple regression will typically show the opposite trend: the more unreliable the covariate, the more the multiple regression actually capitalizes on this unreliability by conflating the direct and indirect effects of the predictor of interest, leading to biased, inconsistent parameter estimates and inflated test statistics. The net effect is that, as the reliability of a covariate falls, it typically becomes easier to reject the null with multiple regression (resulting, as we have already seen, in very high false positive rates when the null is true)...” Jacob Westfall and Tal Yarkoni, ‘Statistically Controlling for Confounding Constructs Is Harder than You Think’, *PLOS ONE* 11, no. 3 (31 March 2016): e0152719, <https://doi.org/10.1371/journal.pone.0152719>.

⁴⁷⁰ Auffhammer, ‘Quantifying Economic Damages from Climate Change’.

⁴⁷¹ Thanks to Danny Bressler for raising this point.

⁴⁷² “The dataset includes a variable measuring regional temperature, originally obtained from the WorldClim database. This variable indicates average temperatures per region between 1950 and 2000.” Christina Greßer, Daniel Meierrieks, and David Stadelmann, ‘The Link between Regional Temperature and Regional Incomes: Econometric Evidence with Sub-National Data’, *Economic Policy* 36, no. 107 (2021): sec. III

Findings of weather studies

Panel studies of weather variation tend to produce higher damage estimates than other types of studies, with one panel study - Burke et al (2015a) - a particular outlier.

Burke et al (2015a) finds that SSP5-RCP8.5 would reduce GDP per capita by 23% by 2100 (median estimate), with substantial uncertainty: there is a 5% chance of a 60% loss in GDP per capita, and in 30% of simulations, they find that climate change increases GDP by 2100.

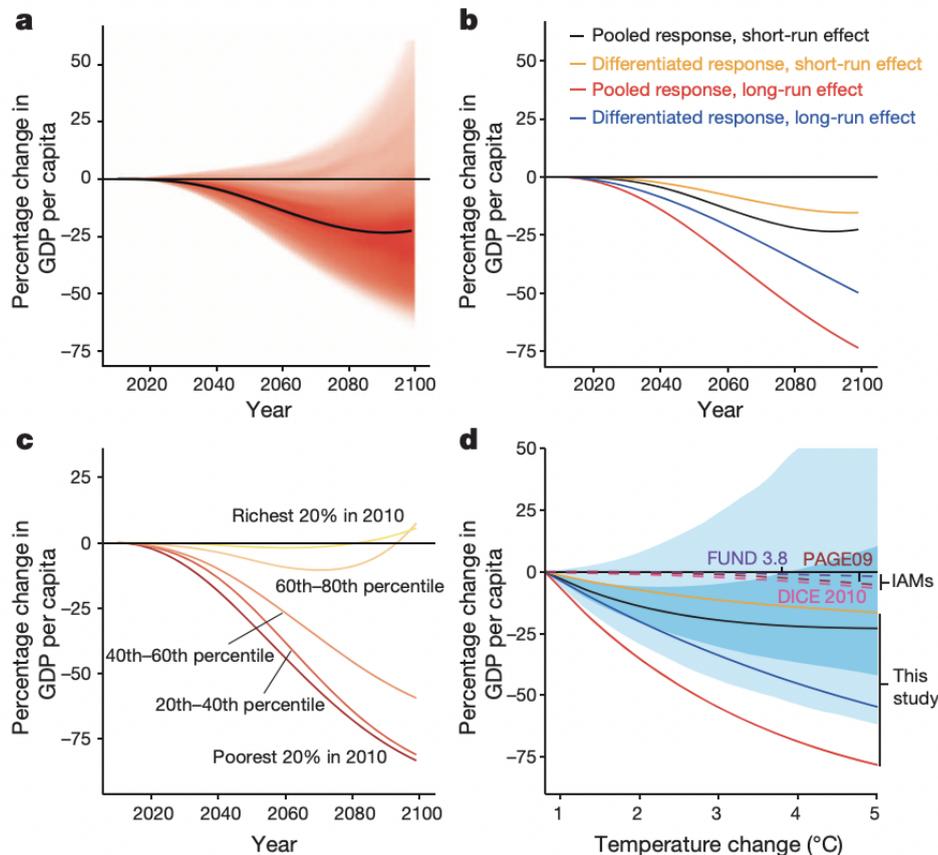


Figure 5 | Global damage estimates arising from non-linear effects of temperature. **a**, Change in global GDP by 2100 using benchmark model (Fig. 2a). Calculation and display are the same as Fig. 4. **b**, Same as **a** (point estimate only) comparing approaches to estimating temperature effects (pooled/differentiated: rich and poor countries assumed to respond identically/differently, respectively; short run/long run: effects account for 1 or 5 years of temperature, respectively; see Supplementary Methods). **c**, Mean impacts by 2010 income quintile (benchmark model). **d**, Projected income loss in 2100 (SSP5) for different levels of global mean temperature increase, relative to pre-industrial temperatures. Solid lines marked as in **b**. Blue shaded areas are interquartile range and 5th–95th percentile estimates. Dashed lines show corresponding damages from major integrated assessment models (IAMs)¹².

Source: Marshall Burke, Solomon M. Hsiang, and Edward Miguel, 'Global Non-Linear Effect of Temperature on Economic Production', *Nature* 527, no. 7577 (November 2015): 235–39, <https://doi.org/10.1038/nature15725>.

Warming has negative effects in both agricultural and non-agricultural sectors and in both rich and poor countries. They find that adaptation so far has been limited.

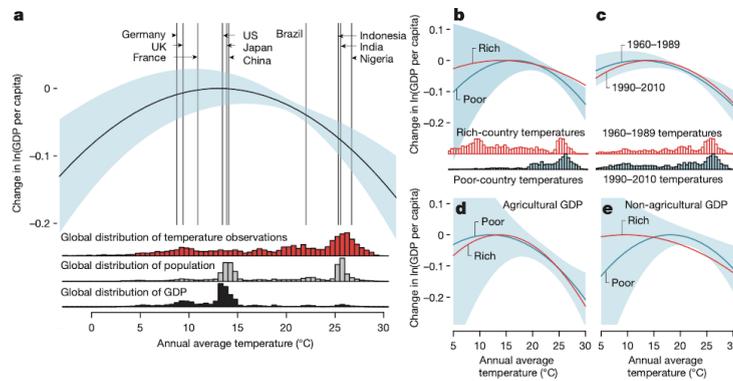


Figure 2 | Effect of annual average temperature on economic production. **a**, Global non-linear relationship between annual average temperature and change in log gross domestic product (GDP) per capita (thick black line, relative to optimum) during 1960–2010 with 90% confidence interval (blue, clustered by country, $N = 6,584$). Model includes country fixed effects, flexible trends, and precipitation controls (see Supplementary Methods). Vertical lines indicate average temperature for selected countries, although averages

are not used in estimation. Histograms show global distribution of temperature exposure (red), population (grey), and income (black). **b**, Comparing rich (above median, red) and poor (below median, blue) countries. Blue shaded region is 90% confidence interval for poor countries. Histograms show distribution of country-year observations. **c**, Same as **b** but for early (1960–1989) and late (1990–2010) subsamples (all countries). **d**, Same as **b** but for agricultural income. **e**, Same as **b** but for non-agricultural income.

They also find highly varying regional effects, with especially bad outcomes in hotter countries. In Sub-Saharan Africa and South and Southeast Asia, the costs of 4-5°C of warming are around 90% of GDP.

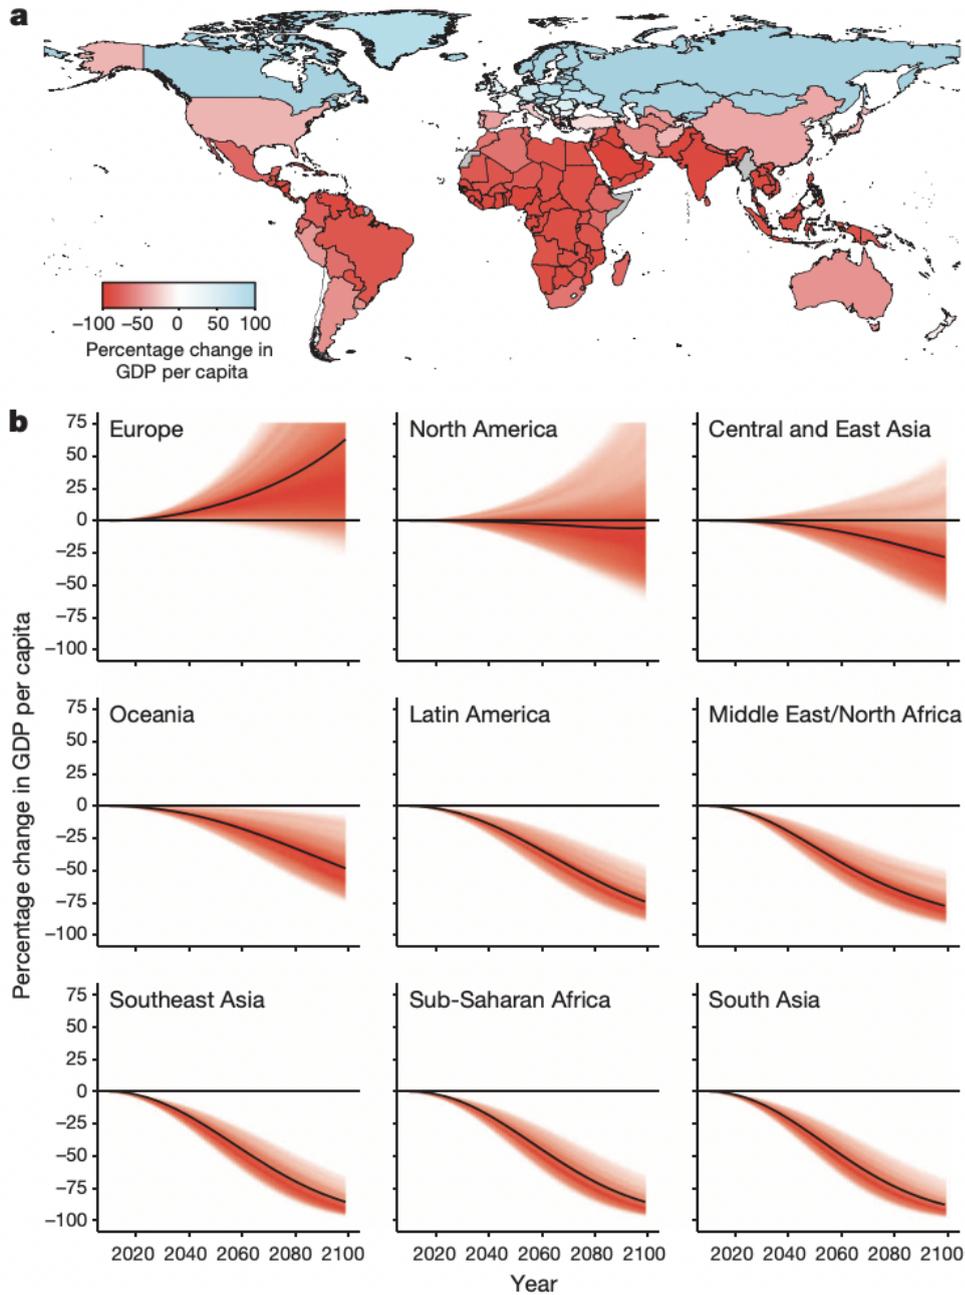


Figure 4 | Projected effect of temperature changes on regional economies. **a, b**, Change in GDP per capita (RCP8.5, SSP5) relative to projection using constant 1980–2010 average temperatures. **a**, Country-level estimates in 2100. **b**, Effects over time for nine regions. Black lines are projections using point estimates. Red shaded area is 95% confidence interval, colour saturation indicates estimated likelihood an income trajectory passes through a value²⁷. Base maps by ESRI.

The finding of Burke et al (2015a) on the country-level costs of climate change has been adopted in the latest version of the PAGE integrated assessment model,⁴⁷³ which has historically been used by the US government to estimate the social cost of carbon.

Other panel studies tend to find costs around 5-15% of GDP for 2-4°C of warming.

Newell et al (2021), which I discuss in more detail below, finds costs to GDP levels of 1-3%, as well as enormous model uncertainty.

GDP growth or GDP levels?

As I discussed above, whether climate change affects the level of GDP or the growth rate is a crucial determinant of the size of climate damages. If temperature shocks merely have a level effect, then after damage in one bad year, the economy will bounce back in the next year. But if they have a growth effect, we would expect to see longer lasting effects of temperature shocks. Thus, we can test for a growth effect by estimating models with lags of temperature.⁴⁷⁴ Both Burke et al (2015) and Dell et al (2012) produce equivocal findings on whether climate change affects the level of GDP or the longer-term growth rate.⁴⁷⁵

It is also important to consider theoretical reasons explaining why climate change would damage economic growth. One important possible mechanism is that level effects in one period could affect the growth in the capital stock, an effect which would be amplified if productive capital were diverted to costly adaptation measures. Another possibility is that climate change affects the speed of technological change, for instance if warmer temperatures affect the cognition needed to produce innovation.⁴⁷⁶ There is scope to ameliorate this latter effect by using air conditioning.

Adaptation

Burke et al (2015a) find that adaptation measures since 1960 have not fundamentally altered the relationship between temperature and productivity. In effect, according to Burke et al (2015a), there was no additional adaptation to the negative economic effects of warmer temperatures between 1960-1989 and 1990-2010, even though average incomes [increased](#)

⁴⁷³ Yumashev notes that PAGE-ICE now includes “new economic impact function based on the recent macro-econometric analysis of the effect of historic temperature shocks on economic growth in multiple countries by Burke et al., projected onto the 8 major regions of the PAGE model using population-weighted temperatures, and adapted to fit with the single year consumption-only approach for climate impacts used in PAGE (also known as level effects).” Dmitry Yumashev, ‘PAGE – ICE Integrated Assessment Models’, 2020, sec. 2.2.

⁴⁷⁴ “Moreover, estimating a model with lags of temperature, they find that this large effect is not reversed once the temperature shock is over, suggesting that temperature is affecting growth rates, not just income levels.22 Growth effects, which compound over time, have potentially firstorder consequences for the scale of economic damages over the longer run, greatly exceeding level effects on income, and are thus an important area for further modeling and research (see Section 4.2).” Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken, ‘What Do We Learn from the Weather? The New Climate-Economy Literature’, *Journal of Economic Literature* 52, no. 3 (2014): 740–98.

⁴⁷⁵ For Dell, see above. For Burke: “Thus, while we can clearly demonstrate that there is a nonlinear effect of temperature on economic production, we cannot reject the hypothesis that this effect is a true growth effects nor can we reject the hypothesis that it is a temporary level effect” Burke, Hsiang, and Miguel, ‘Global Non-Linear Effect of Temperature on Economic Production’, SI p15.

⁴⁷⁶ Burke, Hsiang, and Miguel, ‘Global Non-Linear Effect of Temperature on Economic Production’, SI, p. 10.

by around \$3,000 in this period. The Burke et al (2015) estimate of the damage of climate change is conditioned on the assumption that *future* adaptation to climate change will also be minimal.⁴⁷⁷

In my view, this finding is difficult to believe. Burke et al (2015a) find that warming damages economic performance in rich countries that are mainly reliant on services (and not on agriculture). One of the main causal explanations Burke et al point to that explains this effect is the micro-level evidence on the effects of temperature on labour supply and productivity.⁴⁷⁸ Since this effect can be eliminated by air conditioning, Burke et al (2015a) implies that as people get richer and as they experience higher temperatures, they are no more likely to invest in air conditioning. This is in tension with common sense and with empirical evidence.

Common sense suggests that as people get richer and experience higher temperatures, they will invest in air conditioning. This includes offices and laboratories that will drive technological innovation and economic growth.

Empirical evidence also strongly supports the idea that, other things equal, richer people will buy more air conditioning. Using micro-level data from 16 countries, Davis et al (2021) find that, within countries, air conditioning increases with income

⁴⁷⁷ “If future adaptation mimics past adaptation, unmitigated warming is expected to reshape the global economy by reducing average global incomes roughly 23% by 2100 and widening global income inequality, relative to scenarios without climate change” Burke, Hsiang, and Miguel, ‘Global Non-Linear Effect of Temperature on Economic Production’.

⁴⁷⁸ “Numerous basic productive components of an economy display a highly non-linear relationship with daily or hourly temperature¹. For example, labour supply⁴, labour productivity⁶, and crop yields³ all decline abruptly beyond temperature thresholds located between 20 uC and 30 uC (Fig. 1a–c)” Burke, Hsiang, and Miguel, ‘Global Non-Linear Effect of Temperature on Economic Production’.

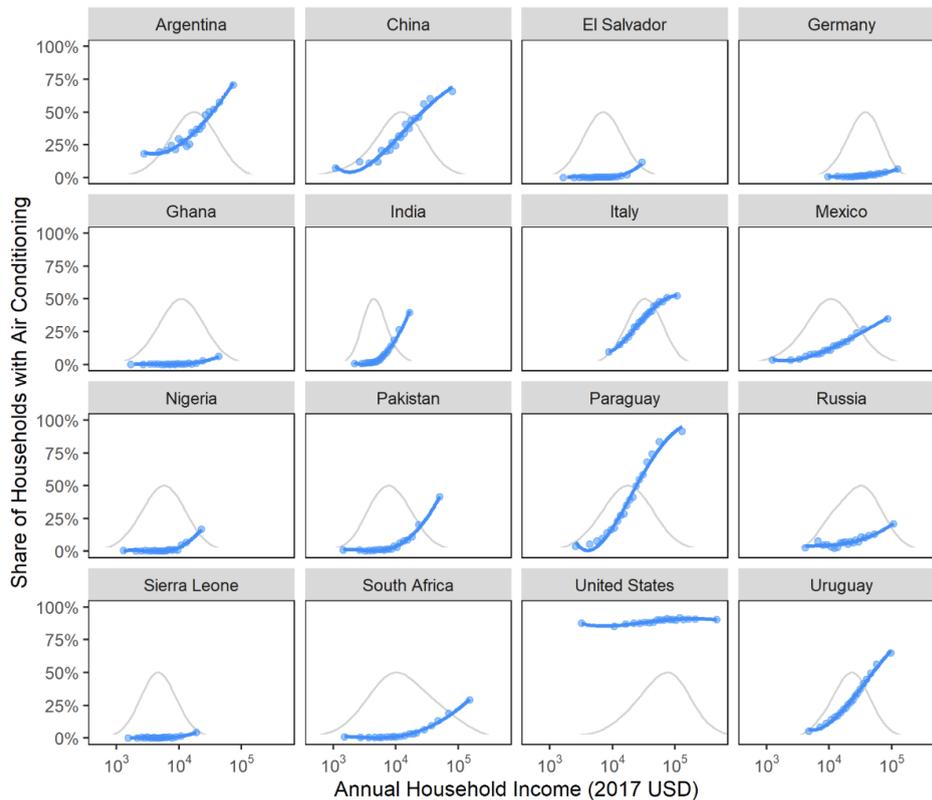


Fig. 1. Air Conditioning Adoption Curves. **Notes:** This figure shows air conditioning adoption curves for each of our sixteen sample countries. The blue points represent the average penetration of air conditioning for each income decile. The blue line is a cubic best fit through these points to illustrate the S-shaped adoption pattern observed in the data. The grey line is a density plot illustrating the income distribution in each country. Income levels throughout have been converted to 2017 US PPP Dollars.

Source: Lucas Davis et al., 'Air Conditioning and Global Inequality', *Global Environmental Change* 69 (1 July 2021): 102299, <https://doi.org/10.1016/j.gloenvcha.2021.102299>.

The global correlation between air conditioning penetration and temperature is confounded by the fact that hotter countries tend to be poorer and so less able to afford air conditioning. However, within countries, there is a strong relationship between temperature and air conditioning penetration. For instance, as of 2008, the coolest regions in the US had penetration rates of 40%, while the warmest and most humid regions were nearly completely saturated.⁴⁷⁹

Since the finding of Burke et al (2015) conflicts with common sense, empirical evidence and with one of their main explanations of how temperature affects economic output, this makes me sceptical that their empirical finding is real.

⁴⁷⁹ "Henderson shows data, including CDD and air conditioner saturation (including both room units and central systems) for the nine U.S. Census Divisions, plus the four largest states (California, Texas, New York and Florida), as provided by the U.S. Energy Information Administration's Residential Energy Consumption Survey (RECS) for 2001. The data show a clear trend, with the coolest regions (Pacific and California) having saturation rates of about 40%, and the warm, humid regions nearly saturated. Henderson references a study which made a fit to U.S. data based on 39 individual cities (Sailor 2003)." Michael A. McNeil and Virginie E. Letschert, 'Modeling diffusion of electrical appliances in the residential sector', *Environmental Energy Technologies Division*, August 2010.

Model uncertainty

One common concern with single studies involving noisy data, limited discipline from theory and complex econometrics is that researcher degrees of freedom and reporting bias combine to make it more likely for researchers to find and/or report negative effects, and for the true extent of model uncertainty not to be presented.

In this vein, Newell et al (2021) argue that the growing literature on weather and the economy makes seemingly *ad hoc* decisions about how to model the relationship between changing climate variables and economic output. Using the same data employed by Burke et al (2015a),⁴⁸⁰ Newell et al (2021) cross-validate different models by estimating models over a subset of the data and then testing how well they perform on the rest of the data. They find enormous model uncertainty:

“Model uncertainty is comparable in magnitude to sampling uncertainty, yielding among GDP growth models a 95% confidence interval for GDP impacts in 2100 of -84% to +359%. GDP levels models yield a much narrower 95% confidence interval of -8.5% to +1.8% and centered around losses of 1-3%, consistent with damage functions of major integrated assessment models.”

If Burke et al (2015a) had used the Dell et al (2012) method of controlling for country-level heterogeneity, their model of non-linear temperature effects would have estimated a 45% *gain* in GDP by 2100.⁴⁸¹

Newell et al's claim that model uncertainty is much higher than Burke and others acknowledge seems intuitively plausible. For example, in Sub-Saharan Africa, and South and Southeast Asia, the 95% confidence interval for damages of 4°C of warming is around 70-90% of GDP. Burke et al (2015a) are very confident that warming will almost completely destroy the economies in these regions, solely on the basis of historical weather data. This is in spite of the fact that such data cannot account well for events usually considered to be catastrophes, such as tipping points or dramatic changes in crop yields or droughts due to long-term warming.

There is some back and forth on this topic between Marshall Burke ([here](#)) and Steve Sexton, a co-author on the Newell et al (2021) paper ([here](#)). The debate centres on whether one thinks cross-validation is a good test of how good a model is at causal inference, or whether one thinks that there are independent reasons to think the controls used in Burke et al (2015a) identify causal effects.

Overall, I am more sympathetic to Newell et al (2021), though I have not looked into Burke et al's arguments for their controls. If the model used in Burke et al (2015a) is a good model of the causal effect of weather on economic performance, then we should expect the model to perform well in cross-validation, in predicting subsets of the same overall dataset on which

⁴⁸⁰ “Models are estimated using the same data employed by BHM.” Richard G. Newell, Brian C. Prest, and Steven E. Sexton, ‘The GDP-Temperature Relationship: Implications for Climate Change Damages’, *Journal of Environmental Economics and Management*, 20 March 2021, 102445, <https://doi.org/10.1016/j.jeem.2021.102445>.

⁴⁸¹ This is pointed out by Steve Sexton, a co-author on Newell et al (2021), [here](#).

the model is based. In his response to Newell et al (2021) Marshall Burke argues that there is a difference between inference and prediction:

“But what if your goal is causal inference? i.e, in our case, in isolating variation in temperature from other correlated factors that might also affect GDP? It's not clear at all that models that perform the best on a prediction task will also yield the right causal results. For instance, prices for hotel rooms tend to be high when occupancy rates are high, but only a foolish hotel owner would raise prices to increase occupancy (h/t Susan Athey who I stole this example from). A good predictive model can get the causal story wrong.”

As Steve Sexton says in his response, since weather is exogenous, it is plausibly causal, whereas hotel prices are obviously endogenous, so this example does not work.

Overall judgement on weather studies

My overall view on weather studies and in particular Burke et al (2015a) is as follows.

Firstly, I am more sceptical of the econometrics in these studies compared to others, for several reasons. There is more scope for researcher discretion when there is complex econometrics and noisy data. My prior expectation is that due to reporting bias, researchers will tend not to disclose the true extent of model uncertainty. Newell et al (2021) provides some confirmation for this.

In addition, there is significant noise and error in the proxies of local temperature and precipitation used in these studies, which makes the risk of false positives much higher. Since the finding of Burke et al (2015) conflicts quite strongly with common sense and the wider impacts literature, I am more inclined to believe that the econometrics is mistaken than that that their estimate is accurate.

Secondly, panel studies miss some particularly important avenues of impact, including:

1. The transition costs of long-term changes in temperature and rainfall (because these are importantly different to weather variation)
2. The costs of tipping points
3. The costs of longer-term impacts such as
 - Sea level rise
 - Ocean acidification
 - CO₂ fertilisation
4. Non-market damages like illness and death

Thirdly, the weather studies try to establish a damage function only using statistics. But we also need a plausible account of the causal mechanism for these damages. In Sub-Saharan Africa, and South and Southeast Asia, the 95% confidence interval for damages of 4°C is around 70-90% of GDP: a huge effect. Moreover, they find this result even though they do not take account of potential tipping points or other avenues of impacts usually considered to be catastrophic.

This finding is strongly at odds with bottom-up studies which try to add up the costs of the most important climate impacts identified in the literature. For instance, Takakura et al (2019) find that the costs of SSP5-RCP8.5 in Asia and Africa are 5% to 15% of GDP. For the Burke et al (2015a) finding to be true, the bottom-up studies must be missing a huge avenue of impact *which can be identified in historical weather data*. This is implausible.

For these reasons, I do not put much weight on the very pessimistic estimate found in Burke et al (2015a). In general, I put more weight on the bottom-up studies than the weather studies.

10.6. Expert elicitation

One criticism of existing economic models of climate change is that they merely embody the judgements of the modeller about the costs of climate change, and are therefore arbitrary. One possible solution to this is to use the opinions of groups of experts for key model inputs.⁴⁸² Thus far, expert elicitation studies have asked experts to give their judgement of the aggregate costs of climate change to GDP. Expert elicitation could in principle also be used for a bottom-up approach by eliciting the views of experts on different sectoral impacts.

10.6.1. Methodological problems with expert elicitation

Expert elicitation studies have several problems. One overarching problem is that experts tend to be poor at prediction. The vast majority of experts perform badly at predicting events a few years in the future, and almost no experts perform well for events more than five years out. In Tetlock's forecasting tournament, discussed in *Expert Political Judgment*, complex models and simple algorithms that extrapolated the past to the future outperformed the best-performing experts.⁴⁸³ Thus, it is not obvious that expert elicitation is superior to formal models in this domain.

⁴⁸² "For an economist, relying on expert opinion might not seem very satisfying. Economists often build models to avoid relying on subjective (expert or otherwise) opinions. But remember that the inputs to IAMs (equations and parameter values) are already the result of expert opinion; in this case the modeler is the "expert." This is especially true when it comes to climate change impacts, where theory and data provide little guidance. Also, we would expect that that different experts will arrive at their opinions in different ways. Some might base their opinions on one or more IAMs, others on their studies of climate change and its impact, and others might combine information from models with other insights. The methods experts use to arrive at their opinions is not a variable of interest; what matters is that the experts are selected based on their established expertise." Robert S. Pindyck, 'The Social Cost of Carbon Revisited', *Journal of Environmental Economics and Management* 94 (2019): 140–60.

⁴⁸³ "Translating the predictions of the crude case-specific extrapolation algorithms, as well as the sophisticated time series forecasting equations, into subjective probability equivalents, we discover that, whereas the best human forecasters were hard-pressed to predict more than 20 percent of the total variability in outcomes (using the DI/VI "omniscience" index in the Technical Appendix), the crude case-specific algorithms could predict 25 percent to 30 percent of the variance and the generalized autoregressive distributed lag models explained on average 47 percent of the variance." "Surveying these scores across regions, time periods, and outcome variables, we find support for one of the strongest debunking predictions: it is impossible to find any domain in which humans clearly outperformed crude extrapolation algorithms, less still sophisticated statistical ones." Tetlock, *Expert Political Judgment*, pp. 53-54.

Related to this, it is not clear how to select the relevant group of experts on a particular topic. Forecasting performance seems like the obvious criterion, but few experts have been tested at forecasting, and when tested, the performance of experts has been relatively poor.⁴⁸⁴

One possible solution to this is to weight expert forecasts according to their performance on short-term calibration questions, which is the approach taken in one influential expert elicitation study on sea level rise by Bamber et al (2019).⁴⁸⁵ Another possible solution is to use teams of superforecasters and climate experts. The climate experts could provide relevant information and the superforecasters could help with expert calibration.

A problem specific to expert elicitation studies of the aggregate economic costs of climate change is that the effects of climate change are highly diverse and wide-ranging. The contributors to damages include sea level rise, agricultural impacts, the effects of temperature on labour supply and productivity, the risk of tipping points, the health costs of heat stress, increased tropical disease, indirect effects such as conflict and so on. Probably the best way to aggregate these effects is to elicit expert opinion on each impact, and then to aggregate the results. In contrast, many studies in this subfield ask experts directly for their view on aggregate economic impacts. In my view, this approach is likely to be less accurate than comprehensive enumerative studies using up to date literature.

10.6.2. Findings of expert elicitation studies

The findings of expert elicitation studies on climate damages are summarised below:(hs p9)

Study	Damage estimate (% of GDP)
Nordhaus (1994) ⁴⁸⁶	3°C: -2% to -4% (median, mean). A range of 0% to -20%. 6°C by 2090: -6% to -10%. Range of -1% to -62%
Schauer (1995) ⁴⁸⁷	2.5°C: -3% to -5% (median, mean) with a variance of 71%.
Howard and Sylvan	3°C: -6% to -10% (median, mean). 10% to 20% chance of a loss of

⁴⁸⁴ “Figures 2.5 and 2.6 bolster another counterintuitive prediction of radical skepticism. Figure 2.5 shows that, collapsing across all judgments, experts on their home turf made neither better calibrated nor more discriminating forecasts than did dilettante trespassers. And Figure 2.6 shows that, at each level along the subjective probability scale from zero to 1.0, expert and dilettante calibration curves were strikingly similar. People who devoted years of arduous study to a topic were as hardpressed as colleagues casually dropping in from other fields to affix realistic probabilities to possible futures.” Tetlock, *Expert Political Judgment*, p. 54.

⁴⁸⁵ Jonathan L. Bamber et al., ‘Ice Sheet Contributions to Future Sea-Level Rise from Structured Expert Judgment’, *Proceedings of the National Academy of Sciences* 116, no. 23 (4 June 2019): 11195–200, <https://doi.org/10.1073/pnas.1817205116>.

⁴⁸⁶ William D. Nordhaus, ‘Expert Opinion on Climatic Change’, *American Scientist* 82, no. 1 (1994): 45–51

⁴⁸⁷ Michael J. Schauer, ‘Estimation of the Greenhouse Gas Externality with Uncertainty’, *Environmental and Resource Economics* 5, no. 1 (1 January 1995): 71–82, <https://doi.org/10.1007/BF00691910>.

(2015 and 2020) ⁴⁸⁸	25% loss of GDP indefinitely
Pindyck (2019) ⁴⁸⁹	Assuming no change in climate policy, by 2066 -10% to -13%. By 2150, around -30%.
Howard and Sylvan (2021) ⁴⁹⁰	For 1.2°C by 2025: 5% chance of an 18% loss by 2025. 3°C: -5% to -9% (median, mean). 5°C: -10% to -16% 7°C: - 20% to -25%

Nordhaus (1994)

Nordhaus interviewed 19 experts on climate change (10 economists, four other social scientists, and five natural scientists), each of whom had a working knowledge of economic statistics.

This survey seems to be subject to numerous problems.⁴⁹¹ Firstly, the survey is small and selection seemed to be arbitrary, as Nordhaus acknowledges.⁴⁹² Secondly, some of the respondents did not work on the question at hand.

“At the other extreme are the "other subdisciplines" of economics (those whose principal concerns lie outside environmental economics); these eight respondents see much less potential for the calamitous outcome— 0.4 percent, or about one-30th of the magnitude estimated by the natural scientists”⁴⁹³

Thirdly, the damage estimates of the natural scientists were 20-30 times higher than the economists, which is evidence of correlated bias in one direction or the other.

Finally, the study asked respondents directly about the aggregate costs of climate change. It seems more promising to decompose this question into a set of smaller subquestions about specific impacts.

Schauer (1995)

⁴⁸⁸ Peter Harrison Howard and Derek Sylvan, ‘Wisdom of the Experts: Using Survey Responses to Address Positive and Normative Uncertainties in Climate-Economic Models’, *Climatic Change* 162, no. 2 (1 September 2020): 213–32, <https://doi.org/10.1007/s10584-020-02771-w>; Peter H. Howard and Derek Sylvan, ‘The Economic Climate: Establishing Expert Consensus on the Economics of Climate Change’ (Institute for Policy Integrity, 2015).

⁴⁸⁹ Pindyck, ‘The Social Cost of Carbon Revisited’.

⁴⁹⁰ Peter Howard and Derek Sylvan, ‘Gauging Economic Consensus on Climate Change’ (New York University School of Law Wilf Hall: Institute for Policy Integrity, 2021).

⁴⁹¹ Steve Keen, ‘The Appallingly Bad Neoclassical Economics of Climate Change’, *Globalizations* 0, no. 0 (1 September 2020): 8ff, <https://doi.org/10.1080/14747731.2020.1807856>.

⁴⁹² “In the end, 22 persons (including the author) were invited to participate, but three did not. Although this selection procedure was arbitrary, it was designed to yield both diversity and informed opinion. The respondents consisted of 10 economists, four other social scientists and five natural scientists and engineers.” Nordhaus, ‘Expert Opinion on Climatic Change’.

⁴⁹³ Nordhaus, ‘Expert Opinion on Climatic Change’, 48.

Schauer (1994) interviewed 14 climate experts. Again, this is a small study and the selection of experts was not systematic and seems somewhat arbitrary, as Schauer acknowledges.⁴⁹⁴

This estimate of the social cost of carbon was produced by combining expert estimates of the value of different parameters. The experts were solicited to estimate parameter values only on parameters on which they had expertise.⁴⁹⁵ However, the damage function itself was not decomposed into distinct individual damages.

Howard and Sylvan (2015 and 2020)

The expert selection procedure for this study is superior to earlier smaller studies. Howard and Sylvan reached out to all those who have published an article related to climate change in a highly ranked, peer-reviewed economics or environmental economics journal since 1994. Of 1,103 experts contacted, 365 responded.

Like the other studies, this suffers from the problem that it does not decompose the damage function into easier-to-estimate components.

Pindyck (2019)

Pindyck contacted 6,833 experts in economics and natural science, and received around 1,000 responses. The questions focused on emissions, damages and discount rate. The study does not decompose estimates of the damage function.

Howard and Sylvan (2021)

Howard and Sylvan invited 2,169 PhD economists to participate in the study, of which 738 participated. These economists have all published an article related to climate change in a leading economics, environmental economics, or development economics journal, and their areas of expertise cover a wide range of issues in climate economics.

The study does not decompose the damage function. Moreover, some of the responses seem poorly calibrated. The respondents estimate a 1% chance of a 45% loss of global GDP in 2025, and a 5% chance of an 18% loss due to climate change.⁴⁹⁶ For context, during the Great Depression, US GDP fell by 25% and Japanese GDP fell by 50% after the end of World War Two. Having read the impacts literature, I find it difficult to see how climate change could cause such massive economic damages in the next three years.

⁴⁹⁴ “The experts interviewed were not chosen by a formally established objective process designed to yield a balanced and comprehensive representation of all scientists. Correcting this shortcoming would improve this work.” Michael J. Schauer, ‘Estimation of the Greenhouse Gas Externality with Uncertainty’, *Environmental and Resource Economics* 5, no. 1 (1 January 1995): 71–82, <https://doi.org/10.1007/BF00691910>.

⁴⁹⁵ “Authoritative experts who have published in the field were chosen so that those interested could readily obtain a sense for their dispositions. Since areas of expertise differed, the experts were not asked to estimate every parameter, but rather only those for which they have an expertise”

⁴⁹⁶ “To address outlier estimates and avoid putting our fingers on the scale, we apply a 95th percent confidence interval trimming methodology, implying a damage range of 0% to -18% in 2025. However, as the upper end of this range implies a permanent catastrophic event with a magnitude akin to the Great Depression occurring in under five years, we also conducted 90th percentile trimming to analyze the median estimate. As the median is relatively stable across various trimming ranges, we focus our analysis primarily on the median estimate after 95th percentile trimming”

10.6.3. Overall judgement on expert elicitation studies

My overall judgement on expert elicitation studies is as follows.

More recent expert elicitation studies seem to be methodologically superior to studies in the 1990s. However, all studies fall short of the state of the art in the emerging science of prediction. In all of the studies discussed above, little attempt is made to improve on the calibration or precision of forecasts, for example with calibration training or weighting according to performance on resolved calibration questions. No incentives are provided for accurate forecasts. None of the studies encourage deliberation between experts, and neither do they encourage respondents to explain their reasoning. None of the studies make use of professional forecasters who have proven track records in prediction.

The questions asked in expert elicitation studies seem unmanageably large, asking experts to make aggregated judgements about all climate impacts without explaining their reasoning.

The advantage that these studies have compared to enumerative and statistical studies is that as things stand in the literature, the estimates are designed to include potential tipping points and catastrophic impacts. The expert elicitation studies tend to find that mean costs are greater than median costs, implying that damages are heavy-tailed and strongly influenced by low probability/high-impact events. However, given the other problems with expert elicitation studies, it is hard to put much weight on their estimates of these impacts. The bottom-up studies provide a more reliable guide than current expert elicitation studies.

10.7. Overall view on climate damages

My overall view on climate damages is as follows.

Firstly, there is little indication from any of the climate economics literature that climate change will do anything close to destroying industrial civilisation or causing civilisational collapse.

Secondly, I put most weight on bottom-up studies that incorporate a wide range of the most important climate impacts, and use up-to-date scientific evidence. Both top-down and bottom-up studies exclude tipping points and indirect effects, but top-down weather studies exclude several other important avenues of impact including:

1. The transition costs of long-term changes in temperature and rainfall
2. The costs of longer-term impacts such as
 - a. Sea level rise
 - b. Ocean acidification
 - c. CO₂ fertilisation
3. Non-market damages like illness and death

In addition, the econometrics of some of the weather studies seems to be questionable.

Thirdly, expert elicitation could in principle be useful, but existing studies are flawed. They ask experts to estimate aggregate economic damages from climate change without explaining their reasoning and without any attempt to improve calibration.

Fourthly, bottom-up studies tend to find that the costs of 4°C are around equivalent to a counterfactual reduction in GDP of around 5%. However, bottom-up studies are likely to understate the *direct* costs of climate change because they:

1. Don't include tipping points
2. Aggregate at regional scales and so don't account for intra-regional inequality
3. Don't model growth effects

Still, even once we account for these effects, it is difficult to see how plausible levels of warming could come close to directly causing civilisational catastrophe.

Fifthly, bottom-up studies also do not capture the potential indirect costs of climate change, which I discuss in the next three chapters.

Finally, one message that emerges from all of the climate economics literature is that people living in poor agrarian countries at low latitudes will be especially badly hit by climate change, even though they have done the least to contribute to the problem. We have strong humanitarian reasons to reduce emissions. We also have strong humanitarian reasons to encourage economic growth in these countries, which will reduce poverty and will make them better able to adapt to climate impacts.

10.8. Climate economics and the long-term

Climate economics is useful because it allows us to arrive at ballpark estimates of the aggregate direct costs of climate change. This is useful for several reasons.

10.8.1. Risk of direct extinction or collapse

Climate economics can tell us how plausible it is that climate change will directly cause extinction or the collapse of industrial civilisation. Climate-economy models try to quantify and aggregate the costs of climate change, and they do not provide much indication that direct extinction or civilisational collapse is on the cards.

10.8.2. Size of indirect risks

Although climate-economy models do not directly try to estimate the direct costs of climate change, they do shed light on the scale of the indirect risks posed by climate change. As a rule, the greater the direct effects something has, the greater we should expect its indirect effects to be. I discuss this in more detail in Chapter 12.

10.8.3. Long-run stagnation

If civilisation stagnates at roughly its current level of technology, this would constitute a huge loss of potential future value. If climate damages are high, this would increase the risk of long-term stagnation. However, semi-endogenous growth models suggest that declining

fertility or progress on AI will be the most important determinant of long-run economic growth. Climate change is weakly levered on these things.

Another argument is that, conditional on technological stagnation, climate change makes the future worse because it slows down economic growth for thousands of years. How important this is depends on whether climate change mainly has a levels effect or a growth effect, which remains deeply uncertain. Much of the effect of climate change will have diminished after 10,000 years, so this is small relative to the expected lifetime of humanity.

10.8.4. The relative geopolitical power of different countries

The relative growth rates of different countries or regions could have important implications for the long-term. Firstly, one potential contributor to the risk of great power conflict is the prospect that a rising power will surpass the current hegemon. I discuss how climate change bears on this possibility in Chapter 12.

Secondly, whichever country becomes dominant this century may be able to lock-in their values in perpetuity. The main determinant of this will be progress in AI, and climate change is a weak lever on AI progress. Climate change also seems a weak determinant of future growth in the two main contender hegemon for this century, the US and China. I discuss this in more detail in the Chapter on conflict.

10.8.5. What is the sign of the effect of economic growth on the long-term?

In my view, the sign of the effect of increased economic growth on the long-term is unclear. Economic growth may shorten the time of perils, but it is also a driver of technological risks that threaten the destruction of civilisation and of potential value lock-in by a global hegemon. It is not clear how these factors trade-off against one another.

There is a difference here between progress at the technological frontier and catch-up growth. Progress at the technological frontier is more dangerous because it has greater scope to uncover dangerous new technologies, whereas catch-up growth has less scope to do this. Still, I think catch-up growth does have some effect on emerging technological risk. For instance, as India grows, it will likely spend more on AI development, on biotechnology research and on nuclear weapons.

11. Displacement and migration

The possibility of large numbers of climate refugees or climate migrants is often raised in climate worst-case scenarios. In particular, climate refugees and migrants are sometimes raised as potential stressors of war.

11.1. Definitions and trends

Migration is voluntary whereas displacement is involuntary. The line between migration and displacement is somewhat fuzzy.

11.1.1. Displacement

Displacement can be driven by conflict and violence, or by weather-related and geophysical disasters (like earthquakes and volcanoes). Conflict includes persecution, violence, human rights violation or events seriously disturbing public order.

The chart below shows the flow of new displaced people between 2011 and 2020, with displacements from conflict and violence in orange, and weather-related and geophysical displacements in blue. The vast majority of disasters are weather-related. As this shows, weather-related displacements have been around 23 million per year over the last decade.

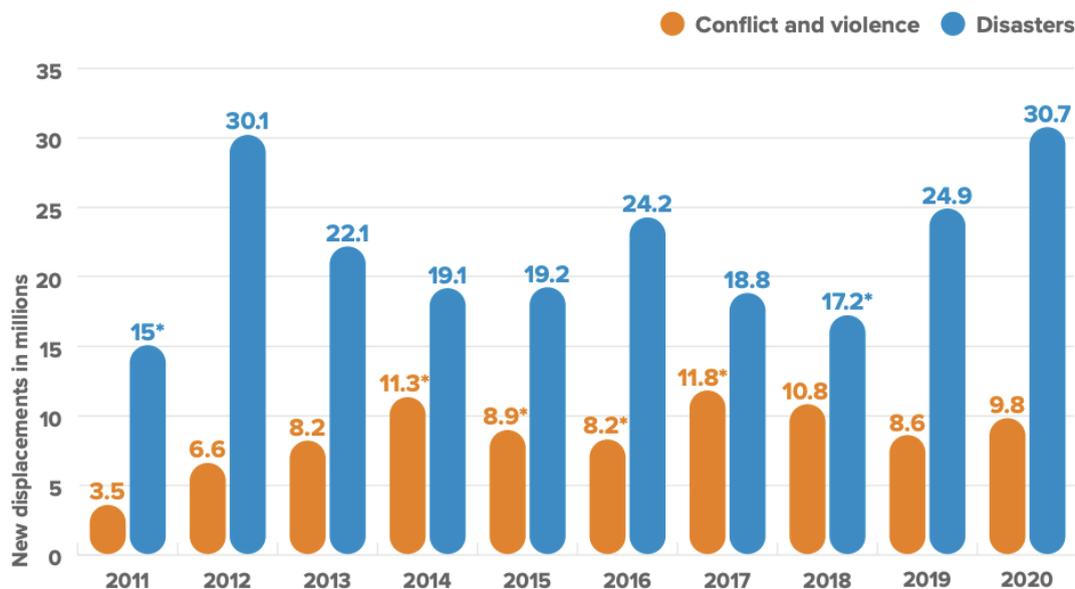


Figure 2: New displacements by conflict, violence and disasters worldwide (2011-2020)

Source: [IDMC, Global Report on Internal Displacement 2021](#)

The chart below shows the breakdown in weather-related displacements by country and type of hazard. Storms and floods account for the vast majority of displacements, with extreme temperatures currently so far responsible for a small minority of displacements. The vast majority of weather-related displacement occurs in Asia. Sub-Saharan Africa, Latin America account for a smaller minority of displacements.

Average annual weather-related displacements, 2010–2020

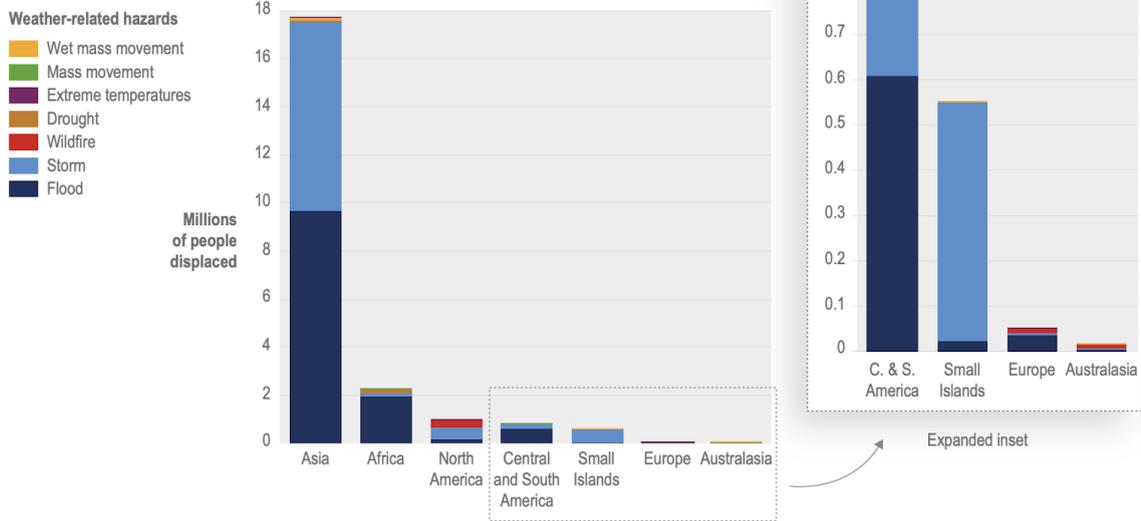
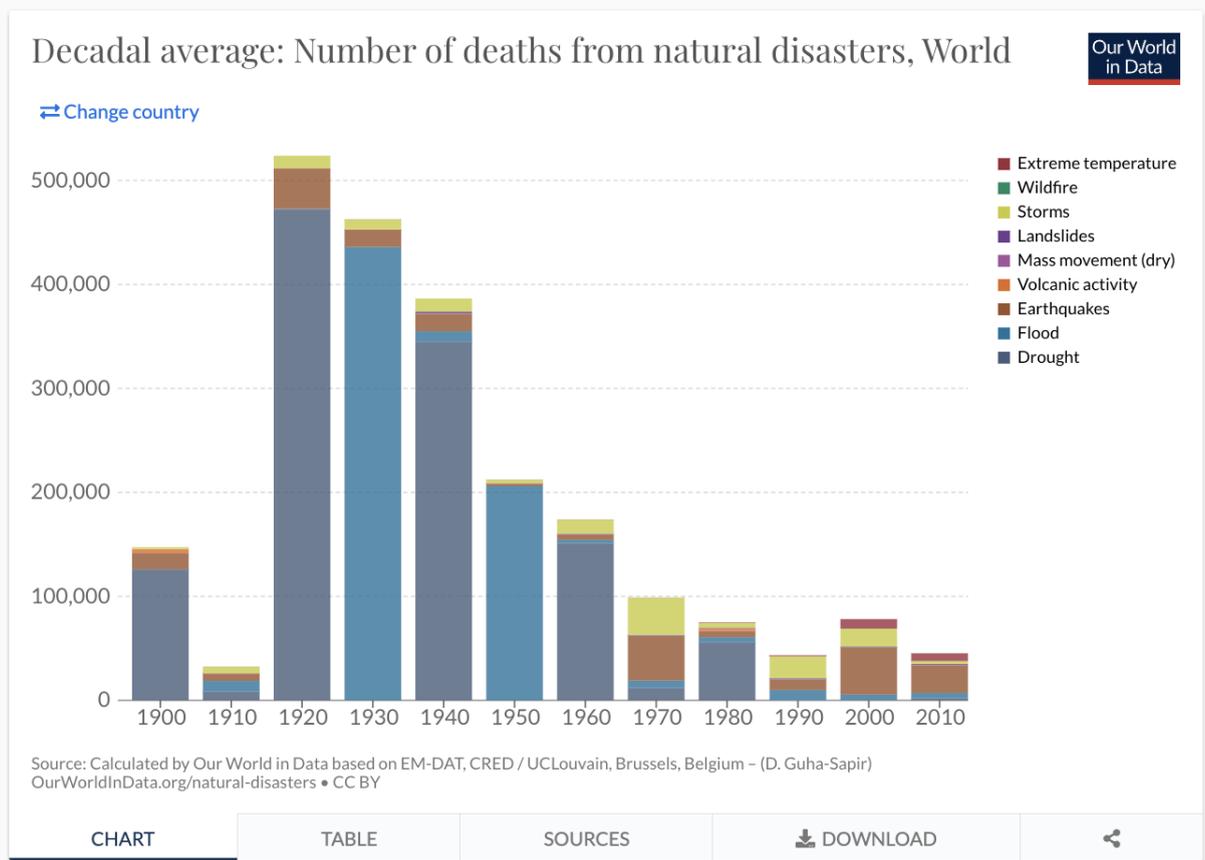


Figure 7.7 | Average number of people displaced annually by selected weather-related events from 2010 to 2020 by region. See text for important notes regarding data collection and trends. Source statistics provided by the Internal Displacement Monitoring Centre (<https://www.internal-displacement.org/>).

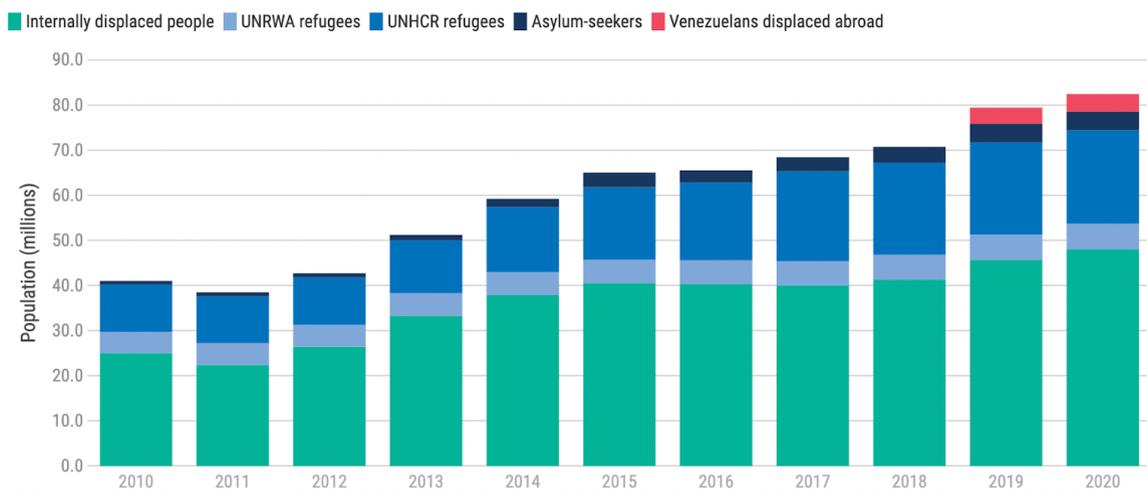
Source: IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022.

I am not aware of longer-term data on weather-related displacements. We do however have data on natural disasters over the 20th Century. Floods and droughts accounted for most natural disaster deaths in the 20th Century, but the absolute number of flood and drought deaths has declined dramatically since the 1920s. Today, earthquakes account for the majority of deaths from natural disasters.



The chart at the start of this Chapter shows the flow of new displacements from conflict or violence. The stock of forcibly displaced people from 2010-2020 is shown below:

Global forced displacement (at end-year)

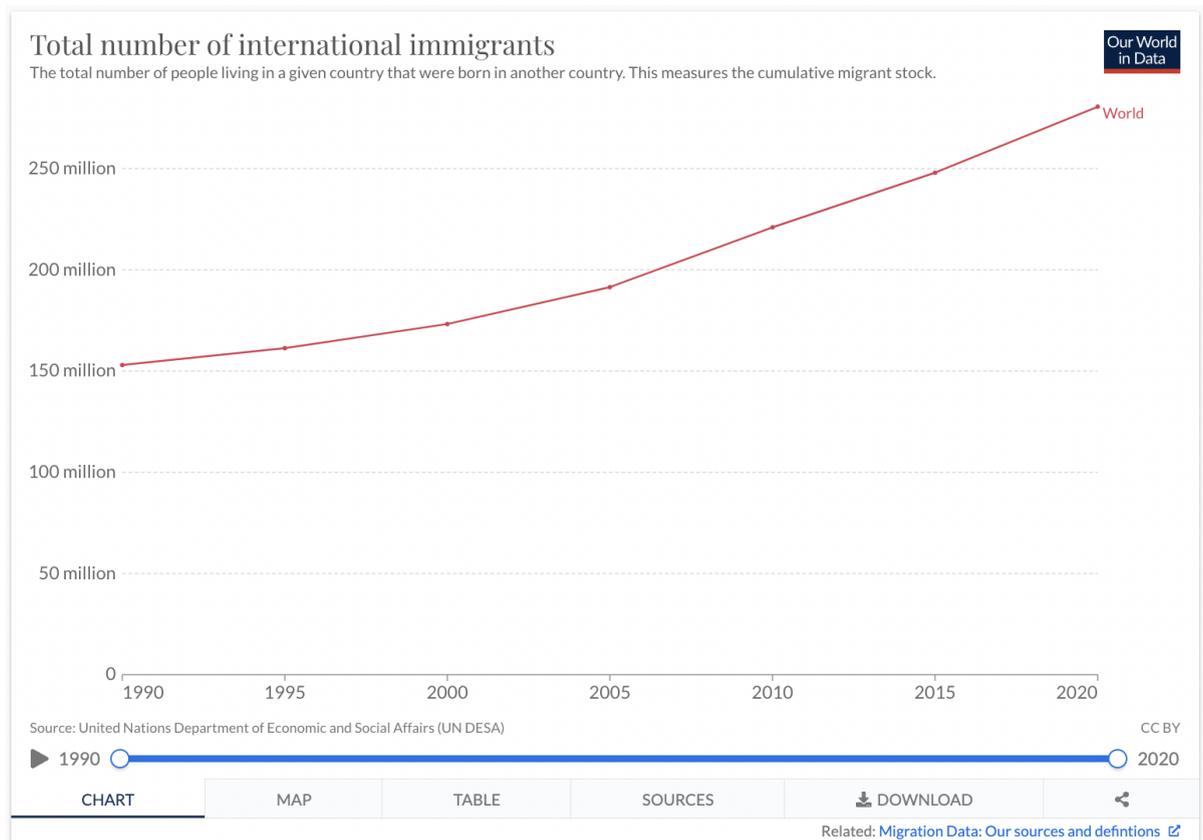


As this shows, most forcibly displaced people remain within their national borders.

11.1.2. Migration

International migrants

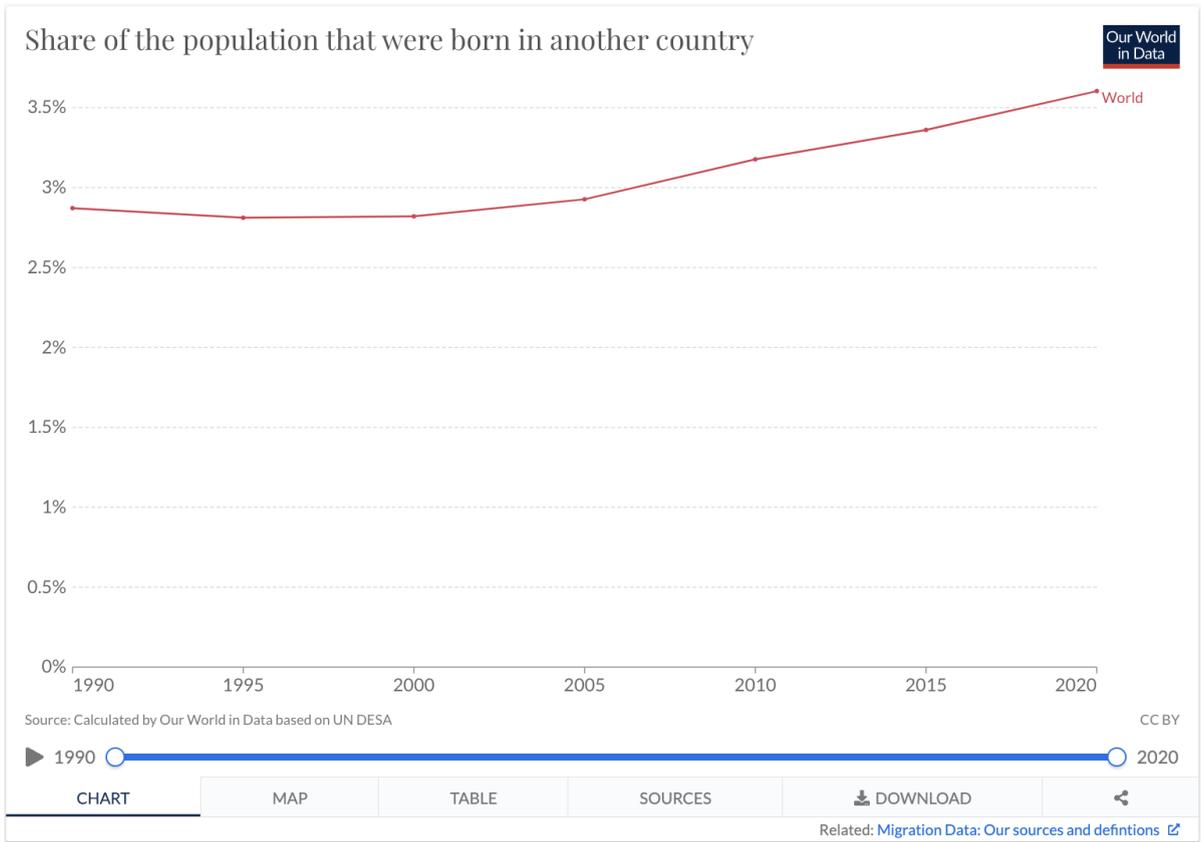
The chart below shows trends in the stock of international immigrants - the number of people in a given country who were born in another country.



This suggests that between 2010 and 2020, the flow of new international migrants was around 6 million per year.

In 1960, around 2% of the world population was born in another country.⁴⁹⁷ This has now increased to 3.66%

⁴⁹⁷ "However, it is clear that international migration stocks have grown over the last 50 years, both in real numbers and as a percentage of the world's population (from 2% in 1960 to 3.1% in 2010)." Foresight p32.



Internal migrants

There were an [estimated](#) 740 million internal migrants in 2009.⁴⁹⁸ I am not aware of any more recent estimates of internal migration, presumably because measurement of internal migration is much harder than international migration.

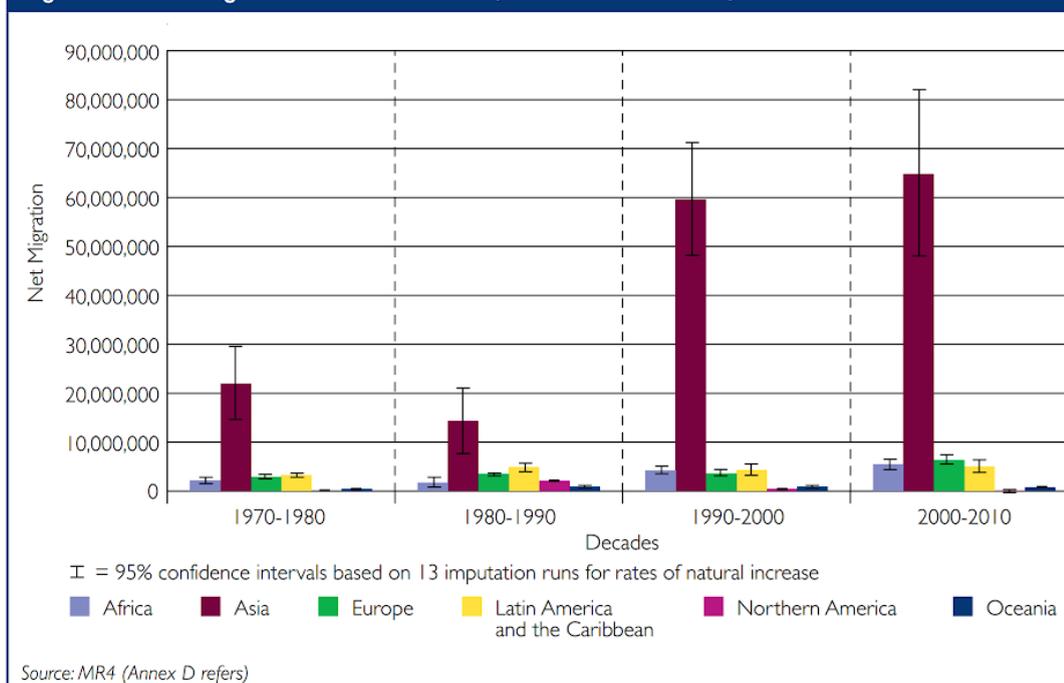
I am also not aware of any estimates of the annual flow of internal migrants. Since the stock of internal migrants is roughly threefold higher than international migrants, we can roughly guess that the flow of internal migrants is around 18 million per year.

Migration to coastal regions

According to the Foresight report for the UK government, between 2000 and 2010, on net, 6.5 million people migrated to coastal regions in Asia per year. This is largely because of superior economic opportunities in coastal regions.

⁴⁹⁸ "The great majority of people do not migrate across borders; much larger numbers migrate within countries (an estimated 740 million internal migrants in 2009).²" International Organization for Migration, 'World Migration Report 2020', p.19.

Figure 3.7: Net migration for coastal areas, within each decade, 1970–2010



Source: Foresight, ‘Migration and Global Environmental Change Future Challenges and Opportunities’ (Government Office for Science, 2011).

11.1.3. The character of migration and displacement

- Environmentally induced population movements are usually internal, temporary and short term.⁴⁹⁹
- Climate-related migration originates most often in rural areas in low- and middle-income countries, and most migrants stay within their own country.⁵⁰⁰
- It is necessary to differentiate between migration caused by slow-onset events, such as droughts and land degradation, and those caused by fast-onset events, such as floods, storms or fires. While the former are usually voluntary and often economically motivated, the latter are involuntary and tend to be short term.⁵⁰¹
- The main motivation for migration is usually economic.⁵⁰²
- Migration is expensive, so within countries, richer groups are more likely to migrate than poorer people.⁵⁰³

11.2. How will climate change affect displacement?

The processes by which climate change will affect migration and displacement are very different. Since migration is voluntary, the main mechanism by which climate change might

⁴⁹⁹ Brzoska and Fröhlich, ‘Climate Change, Migration and Violent Conflict’.

⁵⁰⁰ IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, Ch. 7, sec. 7.2.6.

⁵⁰¹ Brzoska and Fröhlich, ‘Climate Change, Migration and Violent Conflict’.

⁵⁰² Foresight, ‘Migration and Global Environmental Change Future Challenges and Opportunities’ (Government Office for Science, 2011), Ch. 2.

⁵⁰³ Foresight, ‘Migration and Global Environmental Change Future Challenges and Opportunities’ (Government Office for Science, 2011), Ch. 2.

have an effect is by affecting relative economic opportunities in different places. Displacement on the other hand would be caused predominantly by storms and floods.

For this reason, it is much harder to estimate the number of climate change-related migrants than it is to estimate the number of people displaced by climate, though estimating either is very difficult. To estimate the number of climate migrants, we would need to build a model which accurately captures how climate change will affect the economy in all regions. In contrast, the populations of people affected by storms and floods are much narrower (mainly in Asia).

11.2.1. Coastal displacements from storms and floods

In Chapter 7, I discussed the modelling study of coastal displacement by Lincke and Hinkel (2021), which found that 1 metre of sea level rise would increase coastal displacements by around 300,000 people per year relative to 30cm of sea level rise. Displacement will mostly occur in Asia. Their estimates account for the possibility of adaptation by affected countries.

If this estimate is correct, then migration to coastal regions is likely to outweigh displacement from coastal regions. Migration to coastal regions in Asia was about 6.5 million per year between 2000 and 2010, whereas coastal displacement would be in the hundreds of thousands.

11.2.2. River flood displacements

Man Kam et al (2021) model the risk of involuntary displacement due to river flooding on different warming scenarios and population growth. However, they do not account for adaptation by the affected countries,⁵⁰⁴ which is not realistic. So, their estimate can be seen as a very high upper bound; it is not clear what figure is more accurate once we account for adaptation. Since countries in Asia are experiencing relatively fast economic growth, we should expect them to have greater adaptive capacity in the future.

They find that, keeping population constant, for each degree of warming, the global risks of involuntary displacement due to river floods increase by 50%. So, 4°C relative to pre-industrial times would increase displacements from river floods by 150%. Around 10 million people are displaced by floods each year, but I am not sure how these are divided between river flooding and coastal flooding; this data is not available in the latest IDMC report. Floods overall accounted for roughly half of weather-related displacements in 2020 (p11). The IDMC definition of floods includes both coastal and riverine floods.⁵⁰⁵ If we assume that riverine flood displacements account for half of total flood displacements, then around 5 million people are displaced by riverine floods each year. The model of Man Kam et al (2021) implies that 2°C of warming would increase flood displacements by around 2.5 million people per year, while 4°C of warming would increase riverine displacements by 7.5 million per year, without adaptation.

⁵⁰⁴ “On the societal impact side, the population projections do not consider future adaptation measures.” Pui Man Kam et al., ‘Global Warming and Population Change Both Heighten Future Risk of Human Displacement Due to River Floods’, *Environmental Research Letters* 16, no. 4 (March 2021): 8, <https://doi.org/10.1088/1748-9326/abd26c>.

⁵⁰⁵ See the IDMC [codebook](#).

11.2.3. Overall displacements from floods and storms

Overall, it is difficult to know how much climate change will increase displacement in the 21st Century. The Man Tam et al (2021) riverine flood estimates are far higher than the Lincke and Hinkel (2021) coastal flooding estimates, but the former do not account for adaptation, which means that the estimate is likely far too high, though it is difficult to know by how much.

In the absence of better evidence, I would guess that the increase in displacements from floods and storms would be in the low millions per year for 4°C of warming.

11.3. How will climate change affect migration?

There are no reliable estimates of how climate change has affected migration at the global level so far.⁵⁰⁶ According to the IPCC, research so far suggests that weather has mixed effects on migration patterns, and that the effect is context-specific.⁵⁰⁷

Scholars are generally hesitant to make quantitative projections of migration and displacement because forecasting such complex processes is very hard. Nonetheless, there are some quantitative estimates in the literature.

Myers estimates

In several places, Myers (1995) and (2002) claimed that climate change would cause mass migration of hundreds of millions of people. These estimates were used in the Stern Review. These estimates are not taken seriously by the scholarly community.⁵⁰⁸ Myers (2002) makes

⁵⁰⁶ “Reliable global estimates of voluntary climate-related migration within and between countries are not available due to a general absence of concerted efforts to date to collect data of this specific nature, with existing national and global datasets often lacking information on migration causation or motivation.” IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, Ch. 7, sec. 7.2.6.1.

⁵⁰⁷ “A general theme across studies from all regions is that climate-related migration outcomes are diverse (high confidence) and may be manifest as decreases or increases in migration flows, and lead to changes in the timing or duration of migration, and to changes in migration source locations and destinations. Multi country studies of climatic impacts on migration patterns in Africa have found that migration exhibits weak, inconsistent associations with variations in temperatures and precipitation, and that migration responses differ significantly between countries, and between rural and urban areas (Gray and Wise, 2016);(Mueller et al., 2020). Multidirectional findings such as these are also common in single-country studies from multiple regions (A.Call et al., 2017);(Nawrotzki et al., 2017);(Cattaneo et al., 2019);(Kaczan and Orgill-Meyer, 2020). The diversity of potential migration and displacement outcomes reflects (1) the variable nature of climate hazards in terms of their rate of onset, intensity, duration, spatial extent, and severity of damage caused to housing, infrastructure, and livelihoods; and (2) the wide range of social, economic, cultural, political and other non-climatic factors that influence exposure, vulnerability, adaptation options and the contexts in which migration decisions are made (Neumann and Hermans, 2015);(McLeman, 2017);(Barnett and McMichael, 2018);(Cattaneo et al., 2019);(Hoffmann et al., 2020) (high confidence)” IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, Ch. 7, sec. 7.2.6.

⁵⁰⁸ “Overall, Myers’ estimate of the number of ‘environmental refugees’ does not rely on any specific methodology: for each region of the world, the number of internally displaced people is considered. On the basis of these figures, Myers makes an estimate of the proportion that could have been displaced because of environmental disruptions. This estimate is based on reports and observations of environmental degradation in the considered region, but no attention is given to an examination of the linkages between environmental change and migration behaviour. In an essentialist fashion, Myers assumes that all people displaced in an area affected by environmental changes have been

various predictions about drivers of environmental displacement, many of which have turned out to be false, and ran against historical trends.⁵⁰⁹

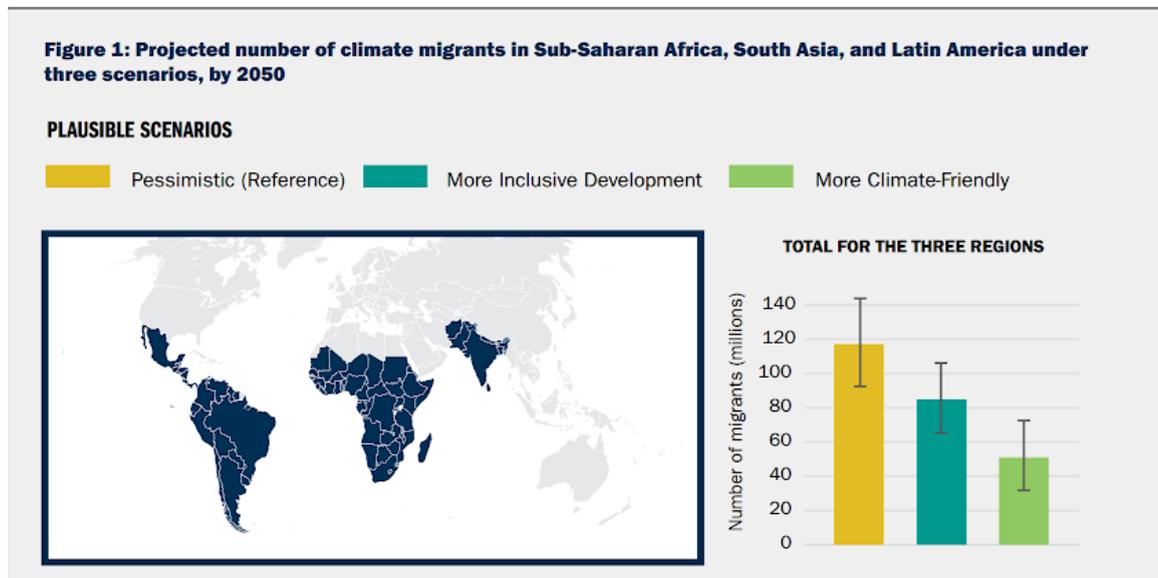
Other estimates

The IPCC's Sixth Assessment Report cites two other sources that try to quantify future migration due to climate change.

Rigaud et al (2018) World Bank report

The first is the 2018 World Bank report *Groundswell* by Rigaud et al.⁵¹⁰ Their model evaluates how climate change will affect water scarcity, crop yields, and sea level rise, and models how this in turn will affect economic push and pull factors in Sub-Saharan Africa, South Asia and Latin America. These regions are among the worst-affected by climate change because they are at low latitude and have a relatively low level of socioeconomic development.

The World Bank estimates that there will be 30 million to 150 million extra internal migrants by 2050, depending on the SSP and RCP.



The scenarios are:

displaced solely because of these changes. Another interesting point to consider is that Myers rules out the possibility that some could have been displaced outside of their country – international migration is not considered in his estimate.” François Gemenne, ‘Why the Numbers Don’t Add up: A Review of Estimates and Predictions of People Displaced by Environmental Changes’, *Global Environmental Change*, Migration and Global Environmental Change – Review of Drivers of Migration, 21 (1 December 2011): S41–49, <https://doi.org/10.1016/j.gloenvcha.2011.09.005>.

⁵⁰⁹ “The total number of malnourished people will continue to grow, with at least 100 million destitutes obliged to live for the most part off imported food.” Norman Myers, ‘Environmental Refugees: A Growing Phenomenon of the 21st Century’, *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 357, no. 1420 (29 April 2002): 609–13, <https://doi.org/10.1098/rstb.2001.0953>. <https://doi.org/10.1098/rstb.2001.0953>. As discussed by *Our World in Data*, undernourishment had in fact steadily declined from 1970 to 2000 and has continued to do so.

⁵¹⁰ Kanta Kumari Rigaud et al., ‘Groundswell’, 2018.

- **Pessimistic:** SSP4, a scenario with high regional inequality, and RCP8.5, which implies about 2.4°C of warming by 2050
- **Inclusive development:** SSP2, the 'trends continue' pathway', and RCP8.5.
- **Climate-friendly:** SSP4 and RCP2.6, which implies 1.7°C of warming by 2050.

However, they do not consider technological improvements in crop yields, which look set to outpace the damage of climate change.⁵¹¹ Since this is one of only three factors that drives migration decisions in their model, I do not put much weight on this study.

The stock of internal climate migrants increases by around 70 million in these regions for 2.4°C relative to 1.7°C of warming between 2020 and 2050. This implies that internal migration increases by around 2.3 million. Given the rough estimate of a flow of 18 million internal migrants per year, this implies that internal migration would increase by around 12%. However, the study also excludes much of the world's population. Overall, the finding could be significantly biased in either direction.

Hoffmann et al (2020)

The IPCC's Sixth Assessment also cites Hoffmann et al (2020) which carried out a meta-analysis of studies exploring the effect of climate change on migration at the country-level. As the chart below shows, studies have produced mixed and uncertain results on the sign of the effect of climate change on migration. The majority of studies find that the mean effect of climate change on *internal* migration is positive, while the findings for international migration are more equivocal.

⁵¹¹ Alex de Sherbinin, personal correspondence, 14 Jan 2020.

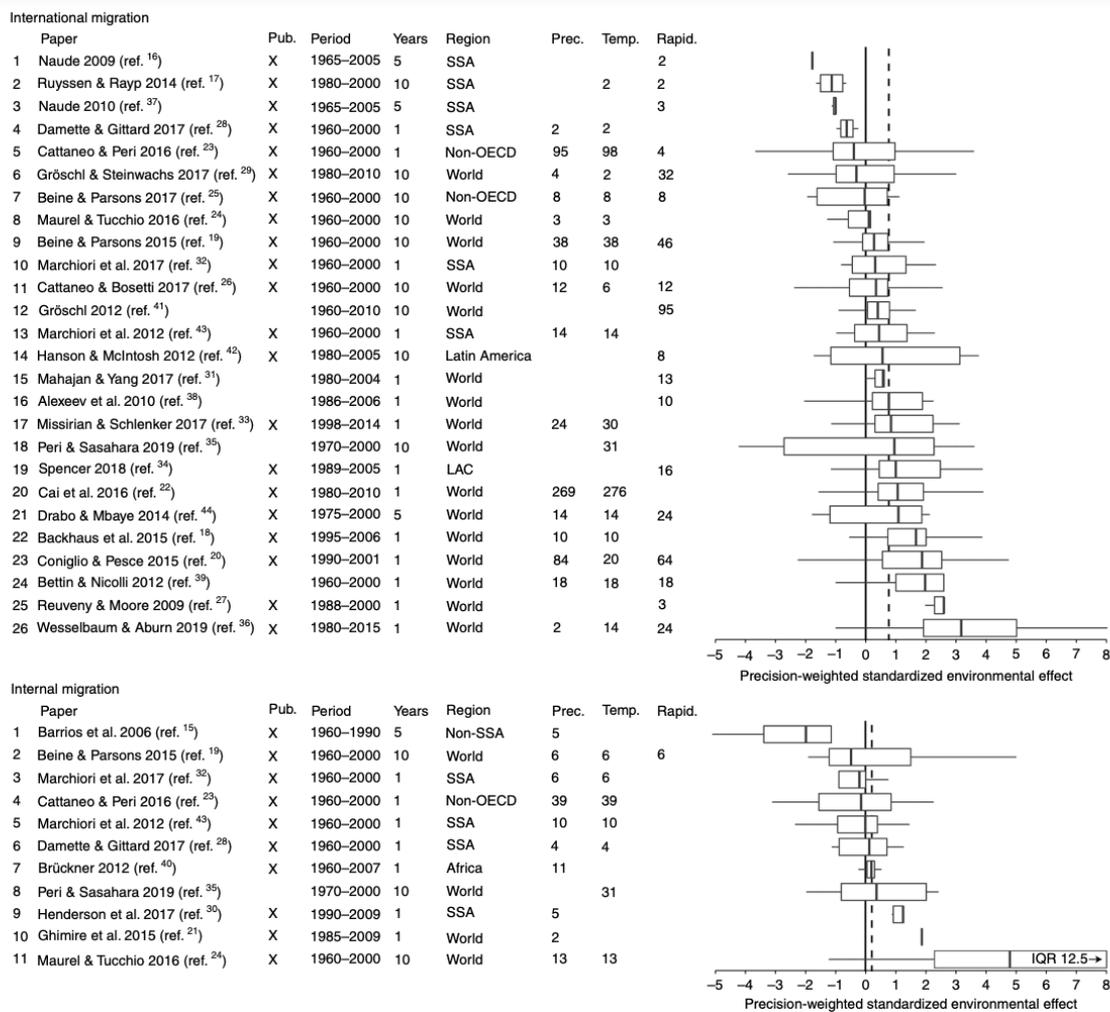


Fig. 1 | Primary quantitative country-level studies testing for a relationship between environmental factors and international and internal migration. 'Pub.' indicates whether the paper was published in a scientific journal listed in SciRef. 'Period' refers to the start and end years of the panel data used in each study. 'Years' captures the most common average period for which the environmental indicators were measured in the study. 'Region' shows the geographical focus of the study. The final columns indicate the environmental factors considered with the figures showing the number of model estimates (*k*) for each factor. Boxplots with median, interquartile ranges (IQR) and whiskers (either maximum value or maximum $1.5 \times \text{IQR}$) of the precision-weighted standardized environmental effects appear on the right. The dashed line shows the median effect size across all studies and estimates. Effects are weighted using the inverse of the estimated variance. 'Prec.' refers to precipitation, 'Temp.' to temperature and 'Rapid.' to rapid-onset disaster events.

Source: Roman Hoffmann et al., 'A Meta-Analysis of Country-Level Studies on Environmental Change and Migration', *Nature Climate Change* 10, no. 10 (October 2020): 904–12, <https://doi.org/10.1038/s41558-020-0898-6>.

The authors produce a funnel plot to test for publication bias, and find no evidence of publication bias.⁵¹²

Hoffmann et al construct different regression models which control for different variables. I discussed the potential issues with omitted variable bias and overcontrolling in Chapter 10. They find that for each standard deviation change in a climate variable, there is 0.021 standard deviation increase in migration.⁵¹³

⁵¹² Roman Hoffmann et al., 'A Meta-Analysis of Country-Level Studies on Environmental Change and Migration', *Nature Climate Change* 10, no. 10 (October 2020): Si, sec. H, <https://doi.org/10.1038/s41558-020-0898-6>.

⁵¹³ "On average, a one standard deviation change in the environmental conditions leads to an increase in migration by 0.021 standard deviations (95%confidence intervals: 0.0176; 0.0235, random effects model, Paule-Mandel estimator, heterogeneity measures $\tau^2=0.002$, $I^2=70.0\%$, $H=1.83$)48"

Hoffman et al also argue that there are various contextual factors that influence migration, including:

- They confirm other migration research which finds that there is a U-shaped relationship between socioeconomic development and migration. Environmental change is smallest among low-income countries, which may be because financial constraints prevent the poorest people from migrating. These financial constraints may be compounded by environmental shocks. Migration rates are highest among middle-income countries.
- The extent to which a country depends on agriculture is a strong predictor of the effect on migration.

They use one of their regression models to project the effects of climate change on migration in different regions.

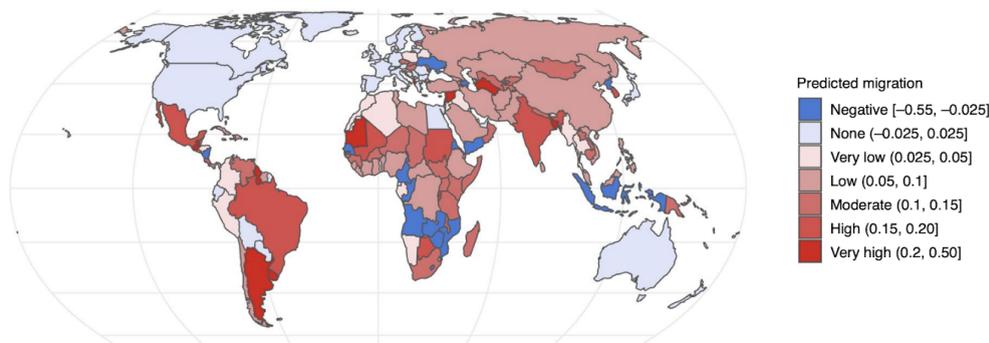


Fig. 4 | Predicted environmental migration worldwide measured in standard deviation changes in migration. Predictions are based on a simplified version of model 5 (Supplementary Table 15), which estimates different migration responses by context. The differential migration responses are combined with information on countries' exposure to environmental change from 1960 to 2000. The derived estimates are integrated into a country-level dataset and used to project the predictions onto the map. The map uses an equal area projection. The term 'negative' refers to a predicted reduction in migration due to environmental change ($-0.55, -0.025$), and 'none' refers to cases with neither a sizeable positive nor negative environmental impact on migration ($-0.025, 0.025$). Positive impacts on migration are subdivided in 'very low' ($0.025, 0.05$), 'low' ($0.05, 0.10$), 'moderate' ($0.10, 0.15$), 'high' ($0.15, 0.20$) and 'very high' ($0.20, 0.50$) impacts. Square brackets in the key denote a closed interval, which includes its limit points. Round brackets denote an open interval not including its limit points.

Hoffmann et al (2020) do not attempt to calculate the overall effect that climate change will have on migration.

Summary of migration estimates

While at first glance it seems plausible that climate change will increase migration and displacement, the overall net effect we should expect is in fact not clear. There are several reasons to think that climate change could reduce migration. For example, by undermining rural livelihoods, it could make costly migration harder for those affected by climate change. At present, there is large net migration to low lying coastal megacities in Asia. If those cities become increasingly affected by climate change, then that may discourage potential migrants.

It is very hard to put a number on overall migration due to climate change.

Roman Hoffmann et al., 'A Meta-Analysis of Country-Level Studies on Environmental Change and Migration', *Nature Climate Change* 10, no. 10 (October 2020): 904–12, <https://doi.org/10.1038/s41558-020-0898-6>.

11.4. Summary of climate change, migration and displacement

My conclusions on migration and displacement are as follows.

Climate-related displacement

The evidence we have suggests that with climate change of 4°C, migration to coastal regions is likely to outweigh displacement from coastal regions. Migration to coastal regions in Asia was about 6.5 million per year between 2000 and 2010, whereas coastal displacement would be in the hundreds of thousands. Migration to coastal regions is likely to increase as population increases and living standards improve.

Estimates of displacements from river flooding are less reliable, but my guess would be that they would be in the low millions for 4°C of warming, predominantly people in Asia. This would increase overall displacements from conflict and weather-related disasters by 5-10%.

Climate-related migration

The literature suggests that climate change will have mixed effects on migration, though most literature suggests that it will increase migration, especially in low- and middle-income countries. It is very difficult to forecast the amount of future migration caused by climate change, and existing estimates are flawed in important ways. Estimates that do exist suggest that climate change would increase the flow of internal migrants by around 10%.

The character of future migration and displacement

Environmentally induced population movements in the future are likely to be similar in character to today: they will usually be internal, temporary and short term. There will likely be a U-shaped relationship between economic development and migration, with the poorest least able to migrate, and middle-income people better able to migrate. Migration caused by sudden onset weather shocks poses fundamentally different challenges to migration caused by slow onset environmental change.

The most important potential effect that displacement and migration might have is increasing political instability and conflict, which I discuss in the next Chapter.

12. Conflict

12.1. Context and trends

12.1.1. Definitions

When thinking about conflict risk, it is important to be careful with different definitions of conflict.⁵¹⁴

- **Interpersonal conflict** - Some studies refer to violent acts that occur between individuals usually described as crimes, such as murder, assault, rape, and robbery.
- **Intergroup conflict** encompasses
 - **Interstate conflict**;
 - **Civil war**, defined by >1,000 battle-related deaths, violence against the government
 - **Civil conflict**, defined by at least 25 battle-related deaths;
 - **Intercommunal violence** (conflict between competing groups within a state)
 - **Low-intensity conflict or social conflict**, e.g., protests and riots
 - **Political repression**.

12.1.2. Trends

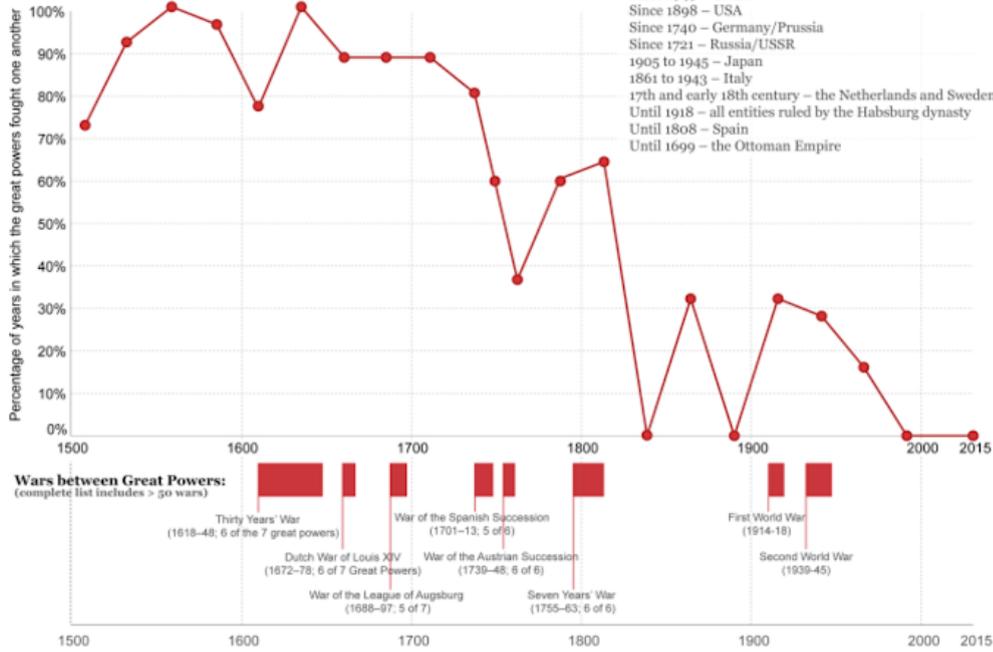
Notwithstanding the recent invasion of Ukraine, conflict of all kinds has declined a lot since the Second World War. Interstate conflict is now very rare. We are in a period known as 'The Long Peace'.

⁵¹⁴ Vally Koubi, 'Climate Change and Conflict', *Annual Review of Political Science* 22, no. 1 (11 May 2019): 343–60, <https://doi.org/10.1146/annurev-polisci-050317-070830>.

Percentage of years in which the 'Great Powers' fought one another, 1500-2015 – by Max Roser

Between 1500 and today there were more than 50 wars between 'Great Powers'. Data are aggregated over 25-year periods.

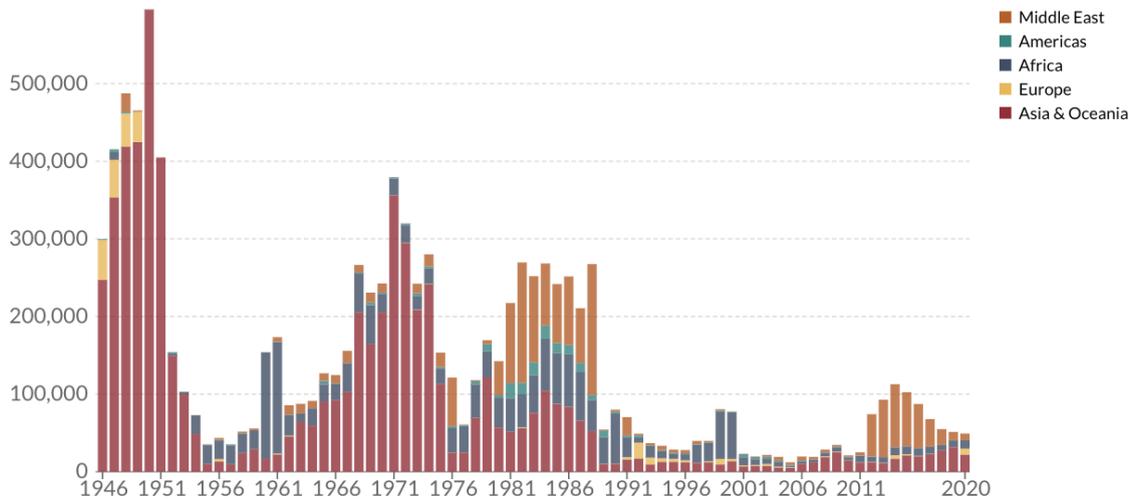
The Great Powers:
 Entire period – France and England/Great Britain/U.K.
 Since 1949 – China
 Since 1898 – USA
 Since 1740 – Germany/Prussia
 Since 1721 – Russia/USSR
 1905 to 1945 – Japan
 1861 to 1943 – Italy
 17th and early 18th century – the Netherlands and Sweden
 Until 1918 – all entities ruled by the Habsburg dynasty
 Until 1808 – Spain
 Until 1699 – the Ottoman Empire



Data source: Steven Pinker (2011) – The Better Angels of Our Nature: Why Violence Has Declined. Based on data from Levy, J. S., & Thompson, W. R. (2011) – The Arc of War. The interactive data visualisation is available at OurWorldInData.org. There you find the raw data and more visualisations on this topic. Licensed under CC-BY-SA by the author Max Roser.

Deaths in state-based conflicts, by world region

Civilian and military deaths in conflicts where the government of a state was a participant on at least one side. The data counts only direct violent deaths (i.e. excluding deaths from disease or famine).



Source: OWID based on PRIO and UCDP. OurWorldInData.org/war-and-peace • CC BY
 Note: The figures shown aggregate the sources' 'best' estimates for deaths in individual conflicts, or the mid-point between high and low estimates where no best estimate is provided.

1946 2020

CHART

TABLE

SOURCES

DOWNLOAD

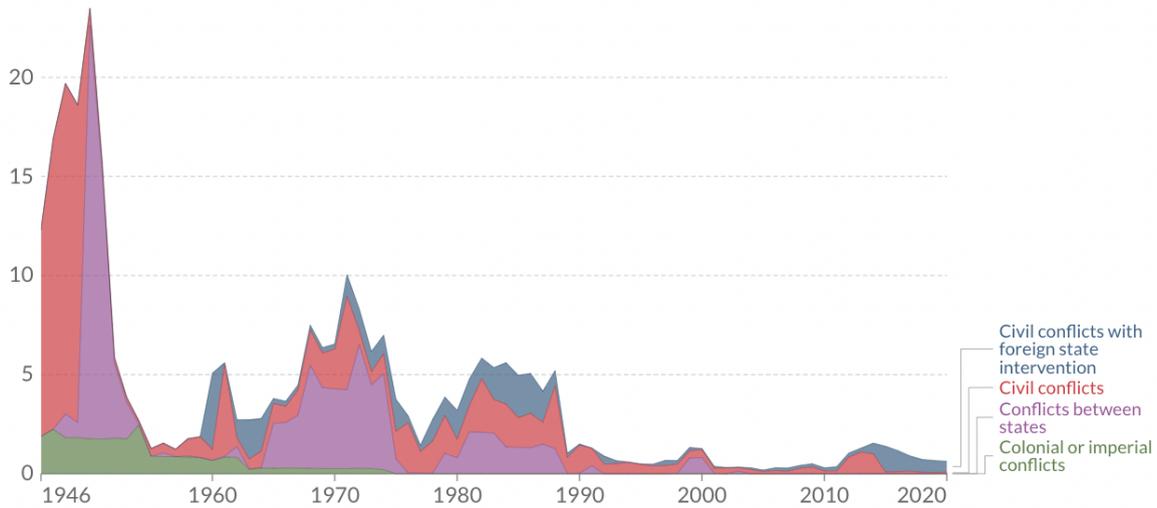
🔗

Related: [Counting world conflict deaths: why do sources differ?](#)

Deaths in state-based conflicts per 100,000, 1946 to 2020

Our World in Data

Civilian and military deaths in conflicts where the government of a state was a participant on at least one side. The data counts only direct violent deaths (i.e. excluding deaths from disease or famine).



Source: OWID based on PRIO and UCDP

OurWorldInData.org/war-and-peace • CC BY

Note: The figures shown aggregate the sources' 'best' estimates for deaths in individual conflicts, or the mid-point between high and low estimates where no best estimate is provided.

▶ 1946 ○ 2020

CHART

TABLE

SOURCES

DOWNLOAD

🔗

Related: [Counting world conflict deaths: why do sources differ?](#)

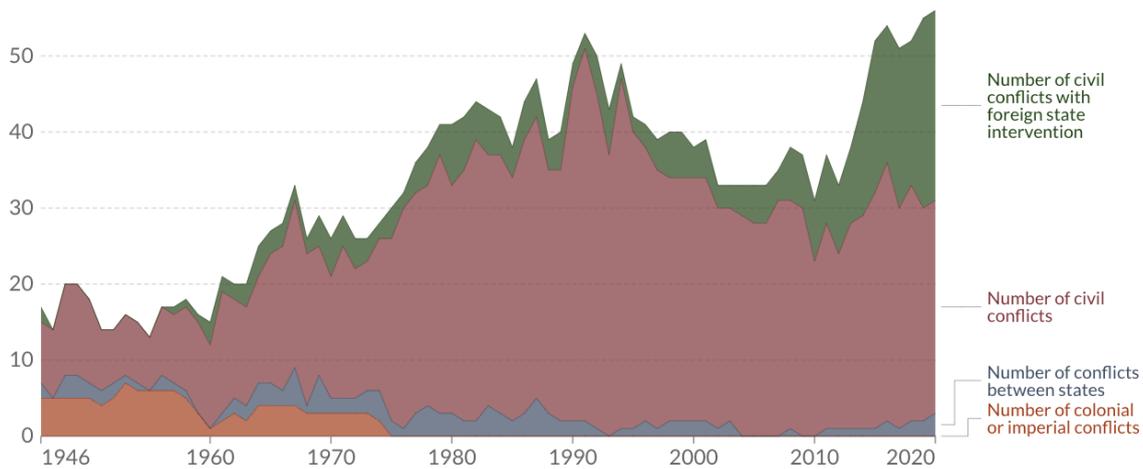
Civil conflicts are now the main type of conflict.

Number of active State-based conflicts, World, 1946 to 2020

Our World in Data

State-based conflicts are conflicts between two parties, where at least one is the government of a state. One-sided violence – such as massacres or genocide – are not included.

↔ Change region



Source: OWID based on UCDP/PRIO

OurWorldInData.org/war-and-peace • CC BY

Note: Ongoing conflicts are represented for every year in which they resulted in at least 25 direct deaths (civilian or military).

▶ 1946 ○ 2020

CHART

TABLE

SOURCES

DOWNLOAD

🔗

Related: [Counting world conflict deaths: why do sources differ?](#)

12.1.3. Drivers of the Long Peace

There are several theories purporting to explain the 'long peace' since 1945, including:

- **Nuclear deterrence.** This is one explanation for the 'long peace' we have observed since 1945.⁵¹⁵
- **Growth, trade and globalisation.** The share of exported goods as a fraction of world GDP has increased substantially since 1945, and some argue that this has helped to reduce the risk of war.⁵¹⁶
- **Democracy.** It is widely accepted in the field of IR that democracies rarely go to war.⁵¹⁷ However, democracies do still go to war with non-democracies.
- **International organisations.** Since World War II, membership of international institutions such as the UN, WTO and EU has increased. Some evidence suggests that countries which share membership in intergovernmental organisations are less likely to go to war with each other, though the strength of this effect is smaller than the effect of trade or democracy.⁵¹⁸

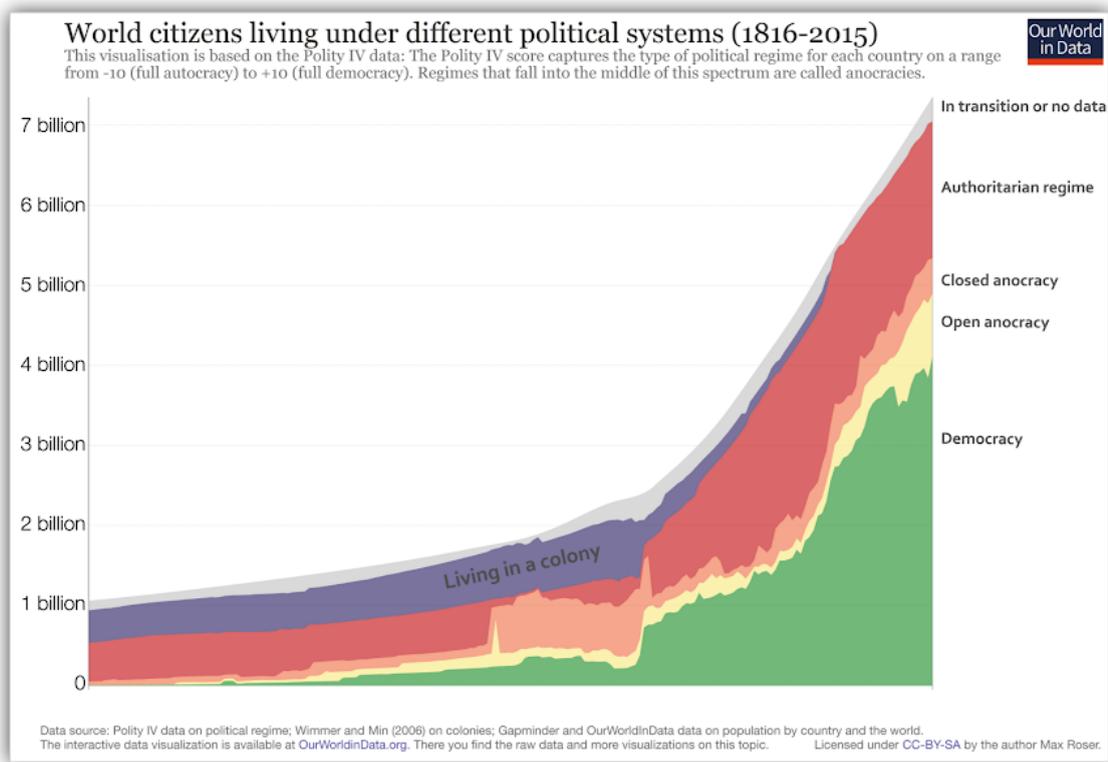
Over the last 200 years, democracy has spread across the globe, as shown on the chart below:

⁵¹⁵ "most [IR researchers] would agree that possession of nuclear weapons vastly reduces the chances of war between two possessors." Greg Cashman, *What Causes War?: An Introduction to Theories of International Conflict*, Second edition (Lanham, Maryland: Rowman & Littlefield Publishers, 2013), p. 362

⁵¹⁶ "increasing economic interdependence from the 10th to the 90th percentile reduces the risk of a fatal dispute by 32 percent." Bruce Russett and John Oneal, 'Causes of Peace: Democracy, Interdependence, and International Organizations, 1885–1992', *International Studies Quarterly* Vol. 47, No. 3 (Sep., 2003), p. 388

⁵¹⁷ "[democratic peace theory has] now gained overwhelming support among Western scholars" (p. 258) Greg Cashman, *What Causes War?: An Introduction to Theories of International Conflict*, Second edition (Lanham, Maryland: Rowman & Littlefield Publishers, 2013), p. 258.

⁵¹⁸ "mutual memberships in intergovernmental organizations (IGOs) make an important independent contribution, above and beyond those of trade and democracy, to reducing militarized disputes between pairs of states (though the pacific benefits of IGO membership are smaller than the other two Kantian variables)" Cited in *What Causes War?*, p. 270



12.2. Climate change and pre-industrial strife

There is research suggesting that pre-industrial cold periods are correlated with war, population decline, dynastic transition and general civilisational strife. The Little Ice Age (1400 to 1900) followed the Mediaeval Warm Period (950 to 1250). In the Little Ice Age, temperatures were often 0.2-0.3°C cooler than they were in the Mediaeval Warm Period. This was correlated with various instances of food price increase, war, epidemics, and population decline.

This chart from Zhang et al (2011) shows this correlation and packs in a lot of information, though it is hard to interpret.⁵¹⁹

⁵¹⁹ David D. Zhang et al., 'Climate Change and Large-Scale Human Population Collapses in the Pre-Industrial Era', *Global Ecology and Biogeography* 20, no. 4 (2011): 520–31, <https://doi.org/10.1111/j.1466-8238.2010.00625.x>.

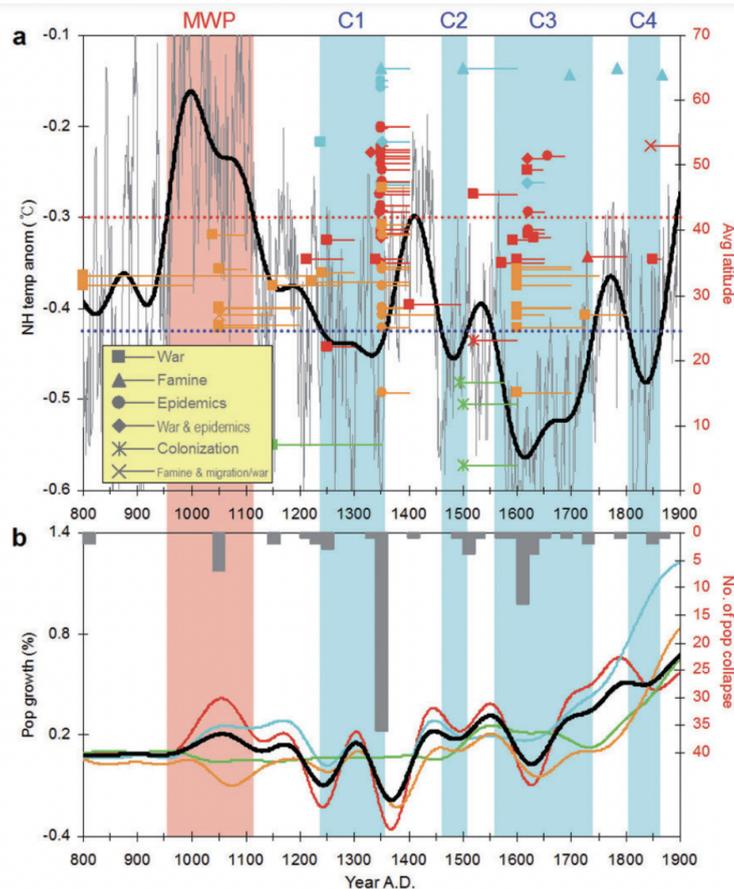


Figure 2 Climate change, population collapse and population growth in the Northern Hemisphere in 800–1900. (a) Northern Hemisphere temperature anomaly (°C) from the 1961–90 mean (Jansen *et al.*, 2007) and population collapse (McEvedy & Jones, 1978). The temperature series (grey line) has been smoothed by the Butterworth 100-year low-pass filter to characterize its centennial variability (bold black line). Red and blue dotted horizontal lines denote the warming threshold (> -0.3 °C) and cooling threshold (< -0.42 °C), respectively. Regarding the details of each population collapse; the average latitude of the relevant country/region corresponds to the right Y-axis and the associated climatic zone is identified by colour (red, warmer humid zone; turquoise, cooler humid zone; bright green, tropical humid zone; orange, dry zone); cause is represented by a symbol (see figure legend); and duration is revealed by the length of the line. (b) Population collapse and population growth rate (McEvedy & Jones, 1978). The grey bars represent the number of population collapses in 20-year units, which corresponds to the right Y-axis (inverted). Population growth rate in different climatic zones is identified by colour (red, warmer humid zone; turquoise, cooler humid zone; bright green, tropical humid zone; orange, dry zone), and bold black (Northern Hemisphere). The red shaded area represents the warm phase in the Medieval Warm Period (MWP), while the blue shaded area represents the cold phases in the Little Ice Age (C1–C4).

Source: David D. Zhang et al., 'Climate Change and Large-Scale Human Population Collapses in the Pre-Industrial Era', *Global Ecology and Biogeography* 20, no. 4 (2011): 520–31, <https://doi.org/10.1111/j.1466-8238.2010.00625.x>.

The red and blue rectangles denote warm and cool periods respectively. The grey rectangles are the size of the population decline: most of them happen in cold periods. The warm period is correlated with variable population effects which are good for certain areas (Europe) but bad for others (Middle East and North Africa). Conversely, the colder period starting around 1550 is *bad* for the Middle East and North Africa.

This chart from a different 2011 paper by Zhang et al shows the correlation between cooling and civilisational strife in the Northern Hemisphere.⁵²⁰

⁵²⁰ David D. Zhang et al., 'The Causality Analysis of Climate Change and Large-Scale Human Crisis', *Proceedings of the National Academy of Sciences* 108, no. 42 (18 October 2011): 17296–301, <https://doi.org/10.1073/pnas.1104268108>.

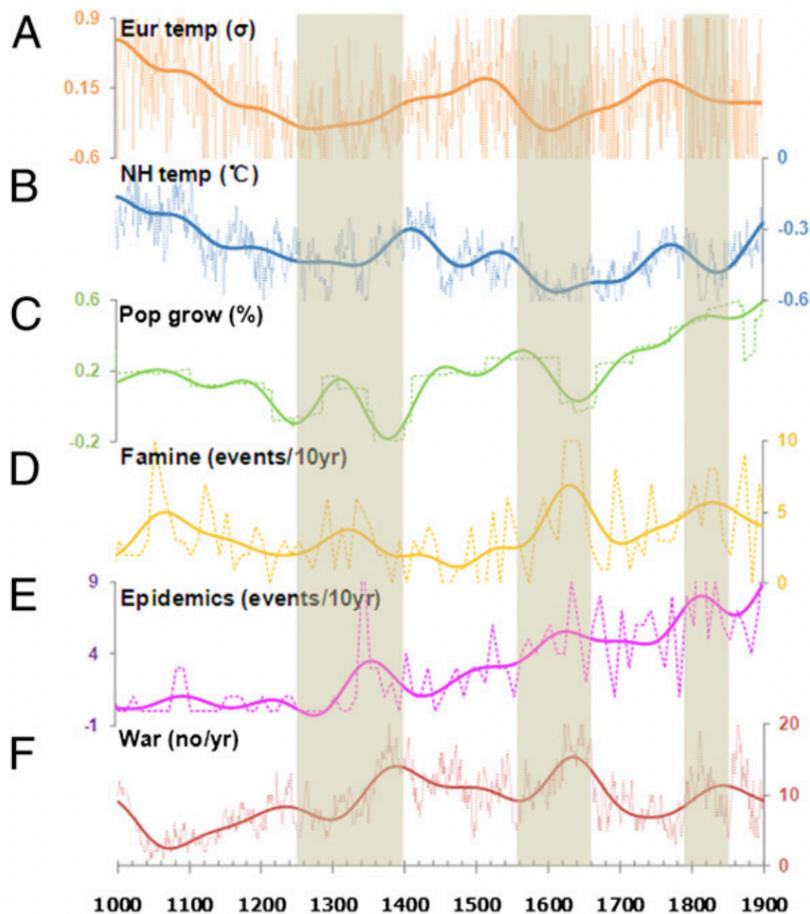


Fig. 4. Temperature change and the alternation of periods of harmony and crisis in the NH during the past millennium. (A) European temperature anomaly (σ). (B) NH temperature anomaly ($^{\circ}\text{C}$). (C) NH annual population growth rate (%). (D) Famine years in the NH (number of famine years per decade). (E) Number of deadly epidemic events (malaria, plague, typhus, measles, smallpox, and dysentery) per decade in the NH. (F) Number of wars per year in the NH. All data were smoothed by a 100-y Butterworth low-pass filter. Gray stripes represent periods of crisis in Europe as delimited by historians (*SI Appendix, Text section 4.2*).

Source: David D. Zhang et al., 'The Causality Analysis of Climate Change and Large-Scale Human Crisis', *Proceedings of the National Academy of Sciences* 108, no. 42 (18 October 2011): 17296–301, <https://doi.org/10.1073/pnas.1104268108>.

Tol and Wagner (2010) find that the correlation between climate and wars declines up until the Industrial Revolution:

“Agriculture became progressively less important over the period, because of economic development, and agriculture became less dependent on the weather, because of improved cultivation methods and better fertilizers.”⁵²¹

⁵²¹ Richard S. J. Tol and Sebastian Wagner, 'Climate Change and Violent Conflict in Europe over the Last Millennium', *Climatic Change* 99, no. 1 (1 March 2010): 65–79, <https://doi.org/10.1007/s10584-009-9659-2>.

This seems plausible. For regions experiencing economic growth, we should expect agriculture to become less dependent on the weather in the future, and for increases in food prices to be less damaging because people are better off.

Moreover, the main effect found in the literature is that in most regions, *cooling* not *warming* causes civilisational strife, so this doesn't provide much evidence that warming will cause civilisational strife.

It is *prima facie* reasonable to be concerned about the methodological soundness of the aforementioned research due to:

- Uncertain climate proxy data
- Uncertain population and other event data
- Researcher degrees of freedom in choosing time periods and classifications of events
- Reporting bias
- Innumerable confounders

12.3. The contemporary relationship between climate change and conflict

12.3.1. The correlation

There is a clear correlation between climate indicators and conflict. In general, countries in the tropics and those suffering water scarcity are more likely to experience civil conflict.

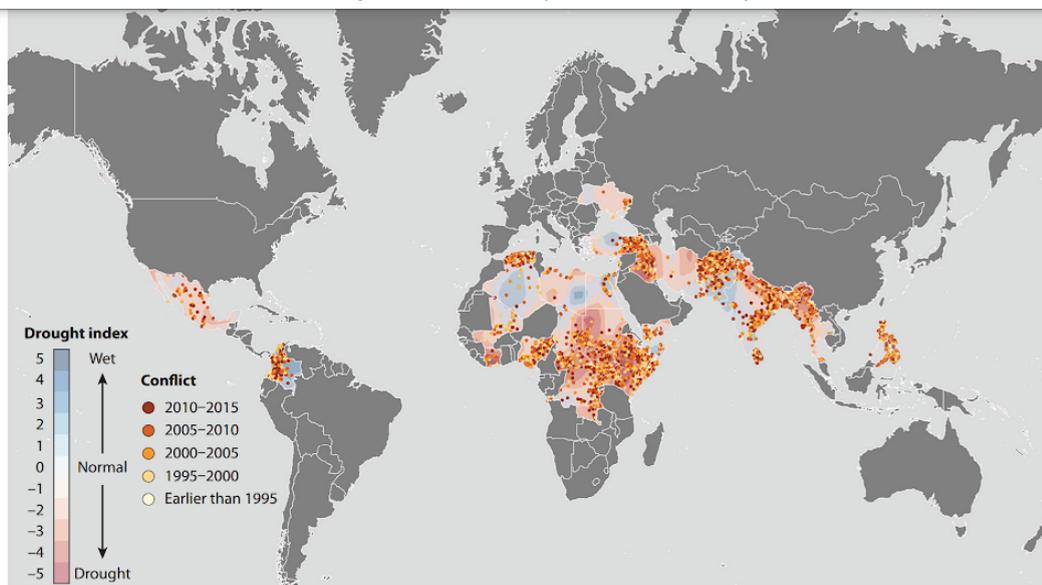


Figure 1

Palmer Drought Severity Index (2005–2014) and location of armed conflict events (1989–2014). Results are screened for countries with more than one recorded armed conflict event per year. Data taken from the National Oceanic and Atmospheric Administration (NOAA PDSI) and Uppsala Conflict Data Program Georeferenced Event Dataset (UCDP GED).

Source: Vally Koubi, 'Climate Change and Conflict', *Annual Review of Political Science* 22, no. 1 (11 May 2019): 343–60, <https://doi.org/10.1146/annurev-polisci-050317-070830>.

Obviously, there are lots of confounders of this relationship. These issues are discussed in more detail in the Chapter on economic costs and institutions.

12.3.2. Overview of the literature

The IPCC's view on the connection between climate change and conflict in the 2022 *Sixth Assessment Report* can be summarised as follows:⁵²²

1. There is as yet little evidence linking climate change to interstate conflict.⁵²³ Consequently, much of the literature focuses on civil conflict.
2. Climate change is widely agreed to be a risk factor for civil conflict and acts through diverse causal mechanisms.⁵²⁴ Changes in agricultural prices are one of the most important mechanisms.⁵²⁵
3. Some studies find that water availability has affected civil conflict, though there is disagreement about this.⁵²⁶

⁵²² This is discussed in sec 7.2.3 and 7.3.3. of IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022..

⁵²³ "Consistent with AR5 findings, there continues to be little observed evidence that climatic variability or 37 change cause violent inter-state conflict." IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3

⁵²⁴ "In intra-state settings, climate change has been associated both with the onset of conflict, particularly in the form of civil unrest or riots in urban settings (high agreement, medium evidence). {Ide, 2020, Multi-method evidence for when and how climate-related disasters contribute to armed conflict risk} as well as with changes in the duration and severity of existing conflicts (Koubi, 2019) Climate change is conceptualised as one of many factors that interact to raise tensions (Boas and Rothe, 2016) through diverse causal mechanisms (Mach et al., 2019);(Ide et al., 2020) and as part of the peace-vulnerability and development nexus (Barnett 2019)(Abrahams, 2020);(Buhaug and von Uexkull, 2021)";

"Potential pathways linking climate and conflict include direct impacts on physiology from heat, or resource scarcity; indirect impacts of climatic variability on economic output, agricultural incomes, higher food prices, increasing migration flows; and the unintended effects of climate mitigation and adaptation policies (Koubi, 2019);(Busby, 2018);(Sawas et al., 2018).Relative deprivation, political exclusion and ethnic fractionalisation and ethnic grievances (Schleussner et al., 2016);(Theisen, 2017) are other key variables. Research shows that factors such as land tenure and competing land uses interacting with market-driven pressures and existing ethnic divisions produce conflict over land resources, rather than a 7scarcity of natural resources caused by climate impacts such as drought. (high agreement, medium evidence) (Theisen, 2017); (Balestri and Maggioni, 2017);(Kuusaana and Bukari, 2015);[also Box 8.3]" IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3.

⁵²⁵ "increases in food price due to reduced agricultural production and global food price shocks are associated with conflict risk and represent a key pathway linking climate variability and conflict (medium confidence) Rises in food prices are associated with civil unrest in urban areas among populations unable to afford or produce their own food, and in rural populations due to changes in availability of agricultural employment with shifting commodity prices (Martin-Shields and Stojetz, 2019). Under such conditions, locally specific grievances, hunger, and social inequalities can initiate or exacerbate conflicts. Food price volatility in general is not associated with violence, but sudden food price hikes have been linked to civil unrest in some circumstances (Bellemare, 2015);(McGuirk and Burke, 2020);(Winne and Peersman, 2019). In urban settings in Kenya, Koren et al (2021) found an association between food and water insecurity that is mutually reinforcing and associated with social unrest (although insecurity in either one on its own was not). Analysing global food riots 2007-2008, and 2011, Heslin (2021) stresses the role of local politics and pre existing grievances in determining whether people mobilise around food insecurity [also Chapter 5]." IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3.

⁵²⁶ "Variation in availability of water has been associated with international political tension and intra-national collective violence (low agreement, medium evidence). Drought conditions have been associated with violence due to impacts on income from agriculture and water and food security, with studies focusing predominantly on sub-Saharan Africa and the Middle East (Ide and Frohlich, 2015);(De Juan, 2015);(Von Uexkull et al., 2016);(Waha et al., 2017);(Abbott et al., 2017);(D'Odorico

4. There is high agreement that climate change has been a small driver of civil conflict relative to other factors, such as socioeconomic development and inter-group inequality.⁵²⁷ Conflict risk is highly mediated by socioeconomic development.⁵²⁸
5. There is disagreement about the size of the effect future climate change will have on conflict.⁵²⁹

Most other reviews of the literature find a lack of scientific consensus and generally conflicting results.⁵³⁰ Sakaguchi et al provide a useful breakdown of neutral, positive and negative findings for the effect of climate change on violent conflict, though, as I discuss below, they do not use the same methodological restrictions as other reviews and meta-analyses:

et al., 2018). A small set of published studies has argued inconclusively over the role of drought in causing the Syrian civil war (Gleick, 2014);(Kelley et al., 2015);(Selby et al., 2017) [also 16.2.3.9]. In general, research stresses the underlying economic, social and political drivers of conflict. For example, research on conflict in the Lake Chad region has demonstrated that the lake drying was only one of many factors including lack of development and infrastructure (Okpara et al., 2016);(Nagarajan et al., 2018);(Tayimlong, 2020). Fewer studies examine the relationship between flooding (excess water) and violence and often rely on migration as the causal factor (see below). However, some studies have shown an association between flooding and political unrest (Ide et al., 2020).” IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3.

⁵²⁷ “Climatic conditions have affected armed conflict within countries, but their influence has been small compared to socio-economic, political and cultural factors (Mach et al., 2019) (high agreement, medium evidence).” IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3.

⁵²⁸ “Future violent conflict risk is highly mediated by socio-economic development trajectories (high confidence). Development trajectories that prioritise economic growth, political rights, and sustainability are associated with lower conflict risk (medium confidence, low evidence).” IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3.

⁵²⁹ “Increases in conflict-related deaths with climate change have been estimated but results are inconclusive (high agreement, medium evidence). Some studies attempted to attribute observed conflict outbreaks to changes in the physical environment and quantify future conflict risk associated with climate change (von Uexkull and Buhaug, 2021);(Theisen, 2017). Burke et al (2015b) concluded that with each one standard deviation increase in temperature, interpersonal conflict increased by 2.4% and intergroup conflict by 11.3%. However, this kind of approach has been criticised for its statistical methods and underrepresenting the known role that socioeconomic conditions and conflict history play in determining the prevalence of violence (Buhaug et al., 2014);(van Weezel, 2019);(Abel et al., 2019).” IPCC, *Climate Change 2022: Impacts*, sec. 7.2.3.

⁵³⁰ Koubi, ‘Climate Change and Conflict’; Kendra Sakaguchi, Anil Varughese, and Graeme Auld, ‘Climate Wars? A Systematic Review of Empirical Analyses on the Links between Climate Change and Violent Conflict’, *International Studies Review* 19, no. 4 (1 December 2017): 622–45, <https://doi.org/10.1093/isr/vix022>.

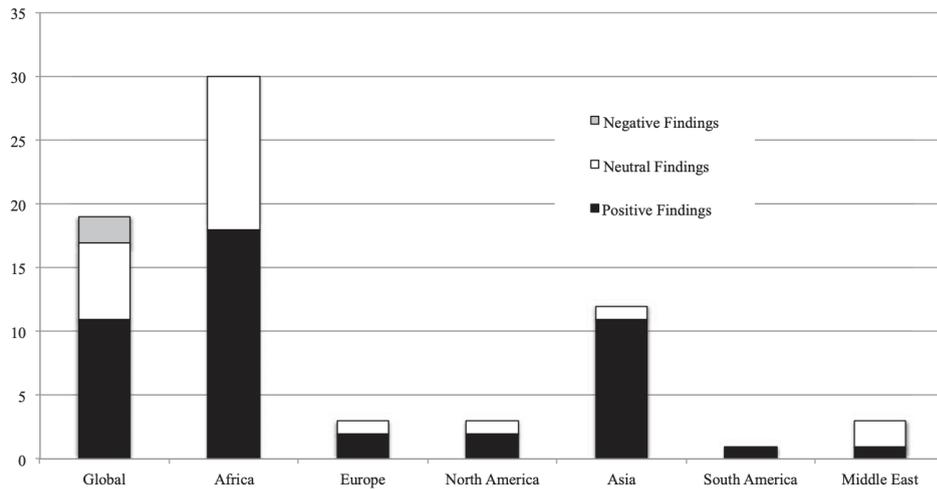


Figure 4. Geographic region of study and study findings.

531

Source: Sakaguchi, Varughese, and Auld, 'Climate Wars?'

They find that a majority of the included studies (62.3%) found small rises in the risk of violent conflict associated with climate change variables. Note that we should expect reporting bias to overinflate the number of positive findings, a point I discuss in more detail below.

However, Hsiang, Burke and others argue that among studies that in their view are not methodologically flawed, there is actually substantial agreement in the literature.⁵³²

This methodological disagreement drives some of the most heated disagreements in the literature, which are between on the one hand, what I will call the 'Berkeley economists' - Burke,⁵³³ Hsiang and others - and on the other hand scholars at the Peace Research Institute at Oslo - Buhaug and others. This is in part driven by a disagreement about the usefulness of cross-sectional regressions and controls for confounders, which I discuss below.

12.3.4. Effect sizes found in the literature

Many of the studies in the literature that find relatively strong consistent effects of climate change are primary research and meta-analyses by the Berkeley economists. Findings for two key studies are shown below:

Study	Study type	Change in relative risk of intergroup conflict per 1-SD increase in temperature	Per 1SD change in precipitation	Interpersonal violence per 1SD change in

⁵³¹ Sakaguchi, Varughese, and Auld, 'Climate Wars?'

⁵³² "We conclude that there is substantially more agreement and generality in the findings of this burgeoning literature than has been recognized previously" Solomon M. Hsiang, Marshall Burke, and Edward Miguel, 'Quantifying the Influence of Climate on Human Conflict', *Science* 341, no. 6151 (13 September 2013), <https://doi.org/10.1126/science.1235367>.

⁵³³ Burke was trained at Berkeley but is a professor at Stanford.

					climate variable
Burke et al (2009), 'Warming increases the risk of civil war in Africa' PNAS	Panel study		27%	NA	NA
Burke et al (2015), 'Climate and Conflict', Annual review of economics	Meta-analysis		11%	3.50%	2.4%

It is not straightforward to translate a standard deviation change into a temperature change. As Burke et al (2015b) note:

“Most studies report changes in climate variables in physical units, such as degrees of temperature or millimeters of rainfall, but different locations around the world exhibit different within-location baseline variances in these measures, which is further exacerbated by differences in the areal extent that is averaged over to compute exposure levels. For example, a 1°C temperature change is a relatively small change for average weekly temperature in a US county; however, it is an enormous change for annual average temperature in an African country. To adjust for these large differences in baseline climate variance, we convert all physical measures of climate into standardized measures based on the within-location standard deviation in climate”

On the A1B emissions scenario, on which there is warming of around 2.5°C by 2050, all regions experience an increase in temperature of 2-4 standard deviations. So, it would be reasonable to assume that for 5°C of warming by 2100, regions would experience warming of 4-8 standard deviations.

These claims about intergroup conflict are highly controversial,⁵³⁴ and have led to heated debate between the Berkeley economists and the PRIO researchers.

Here, I will carry out a '[minimal trust](#)' investigation of the two studies mentioned above: I will suspend my trust in the field and dig as deeply into the question as I can.

12.3.5. A review of Burke et al (2009)

Burke et al (2009) has been cited more than 1,000 times and argues that >1,000 death civil conflict will increase in Africa by 4.5 percentage points per degree of warming, and that this effect will outweigh the effects of growth and democratisation. This is a panel study which examines the effects of interannual weather variation on conflict risk. Heuristically, an economy observed during a cool year is the 'control' for that same society observed during a warmer 'treatment' year. Because interannual weather change is exogenous, any difference in conflict risk is plausibly due to differences in weather rather than country-specific or time trend factors. This methodological approach avoids the omitted variables and

⁵³⁴ H. Buhaug et al., 'One Effect to Rule Them All? A Comment on Climate and Conflict', *Climatic Change* 127, no. 3 (1 December 2014): 391–97, <https://doi.org/10.1007/s10584-014-1266-1>; Halvard Buhaug, 'Climate Not to Blame for African Civil Wars', *Proceedings of the National Academy of Sciences* 107, no. 38 (21 September 2010): 16477–82, <https://doi.org/10.1073/pnas.1005739107>.

over-controlling problems associated with cross-sectional approaches or that try to control for confounders.

As discussed in the Chapter on economic costs, this panel study approach relies on the assumption that interannual weather variation is relevantly similar to long-term climate change, which is not obvious.

I have spent a significant amount of time focusing on this study because much of the posited relationship between climate change and conflict depends on studies of Africa, a point I discuss in more detail below.

There has been vigorous debate about this study with the following critiques and responses

- Buhaug (2010) - '[Climate not to blame for African civil wars](#)', PNAS
- Burke et al (2010) - '[Climate robustly linked to African civil war](#)', PNAS
- Burke et al (2010) - '[Climate and civil war: is the relationship robust](#)', working paper (companion piece to the above)
- Buhaug (2010) - '[Reply to Burke et al.: Bias and climate war research](#)' PNAS
- Buhaug et al (2010) - '[Sensitivity analysis of climate variability and civil war](#)', working paper (expands on the claims in the above Buhaug paper)
- Hsiang and Meng (2014) - '[Reconciling disagreement over climate–conflict results in Africa](#)', PNAS
- Buhaug (2014) - '[Concealing agreements over climate–conflict results](#)', PNAS

My overall judgements are as follows

1. The PRIO researchers do not provide a convincing response to the criticisms of their econometric approach by the Berkeley economists.
2. The Burke et al result is highly sensitive to the choice of time period and becomes much weaker or disappears after 2002, in a way that is inconsistent with one of Burke et al's projections.
3. The Burke et al result is highly sensitive to the definition of civil war.
4. The Burke et al projection is one about severity, not onset, and the severity of conflicts is noisy and declining on the most plausible measures.

Failure to predict decline in civil conflict incidence

Burke et al argue that civil conflict will decline even on an optimistic scenario for growth and democratisation. Specifically, they consider

“an “optimistic scenario,” in which the annual per capita economic growth rate is 2% and the increase in democracy is the same as during 1981–2002, a period of substantial democratic reform in Africa (see Methods).”⁵³⁵

Even on this optimistic scenario, there is little effect on their mean conflict projections

⁵³⁵ Marshall B. Burke et al., 'Warming Increases the Risk of Civil War in Africa', *Proceedings of the National Academy of Sciences* 106, no. 49 (8 December 2009): 20670–74, <https://doi.org/10.1073/pnas.0907998106>.

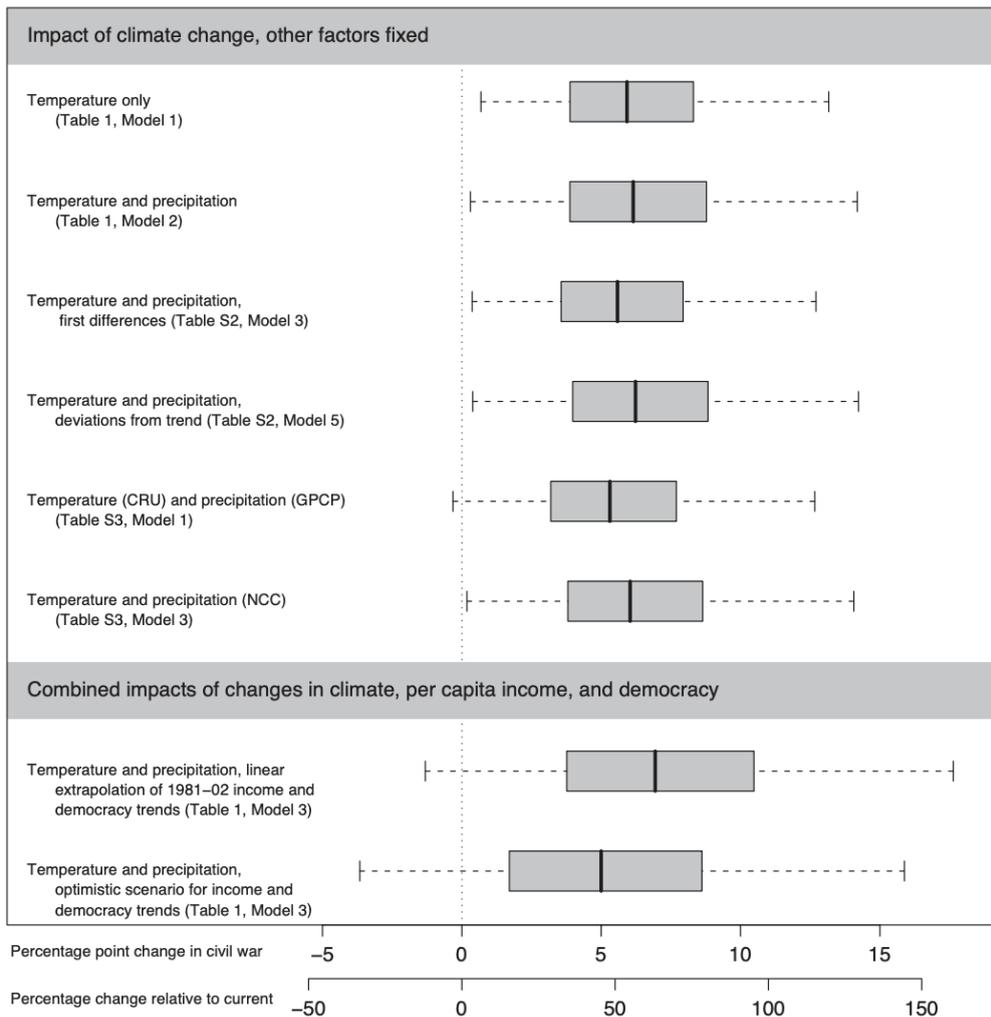


Fig. 2. Projected percent changes in the incidence of civil war for all of sub-Saharan Africa, including both climate and conflict uncertainty as calculated as in Fig. 1. (Top) Projections based on alternative specifications of the relationship between climate and conflict, with other factors fixed. (Bottom) Projected combined effects of changes in climate, per capita income, and democracy. Dark vertical lines represent the median projection, colored boxes show the interquartile range, and whiskers indicate the 5th–95th percentile of projections, using climate projections from all climate models for the A1B scenario, such that each boxplot represents 180,000 projections. Each specification includes the variables listed on the left (contemporaneous and lagged for the climate variables) in addition to country time trends and country fixed effects.

The optimistic scenario is depicted in the bottom box and whisker diagram. They argue that even on the optimistic scenario, climate change would lead to increased war by ~50% by 2030. If we assume that the increase is linear, the increase would be 18% by 2010 and 36% by 2020. The bottom fifth percentile is a 40% decline by 2030. As they note:

“Furthermore, the adverse impact of warming on conflict by 2030 appears likely to outweigh any potentially offsetting effects of strong economic growth and continued democratization.”⁵³⁶

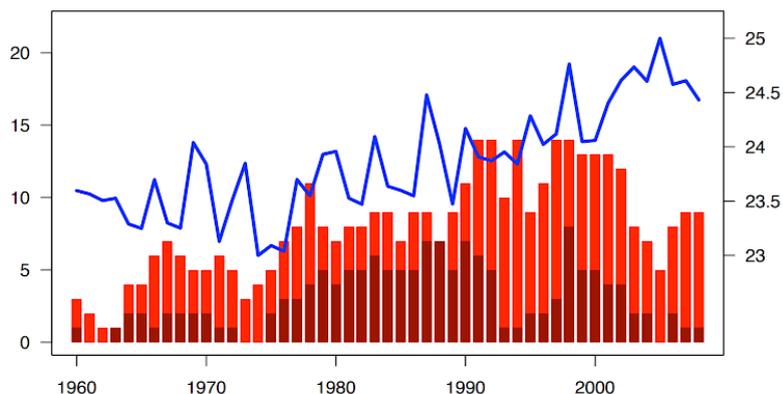
The Burke et al (2009) study only uses data on conflict up to 2002.

However, as Buhaug (2010) points out, between 2002 and 2008, >1000 death civil conflicts declined by about 50% in Africa by 2008, despite warming.⁵³⁷

⁵³⁶ Burke et al.

⁵³⁷ Buhaug, ‘Climate Not to Blame for African Civil Wars’; Halvard Buhaug, ‘Reply to Burke et al.: Bias and Climate War Research’, *Proceedings of the National Academy of Sciences* 107, no. 51 (2010): E186–87.

Figure 1: Trend in Sub-Saharan African conflict and temperature, 1960-2008. Dark red indicates number of large wars (>1000 deaths) ongoing in a particular year, and light red the number of small wars (>25 deaths) ongoing, based on the Uppsala/PRIO data. Blue line indicates continental average temperature (right Y-axis, from UDel data).



538

The light blue is state-based conflicts with >25 deaths. These have not increased by a third by 2012, which is what one might expect on the Burke et al model.

The chart below presents the trend up to 2019 for wars (>1,000 battle-related deaths) and conflicts (25–999 battle-related deaths)

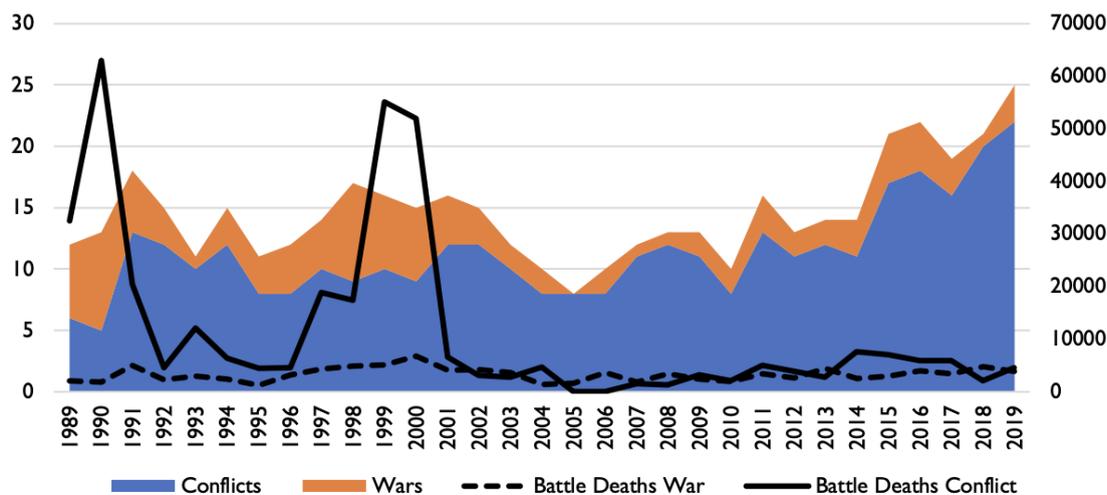


Figure 3: Conflict and battle deaths split between conflicts and wars in Africa, 1989–2019

Source: UCDP/PRIO Armed Conflict Database and UCDP Battle Death Database (Pettersson & Öberg 2020)

Source: Palik et al, '[Conflict Trends in Africa, 1989–2019](#)', PRIO paper, 2020

⁵³⁸ Marshall Burke et al., 'Climate and Civil War: Is the Relationship Robust?' (National Bureau of Economic Research, 2010).

As this shows, the trend in >1000 death civil wars was low and stable between 2002 and 2019.

Applying the Burke et al model to the 2003-2008 period, the relationship between temperature and war now disappears. This is in model 2 in table 4 below

Table 4: Results using UDel climate data, for periods 1980-2002 (Model 1) and 2003-2008 (Model 2). All models include country fixed effects and country-specific time trends, with Huber-White standard errors adjusted for clustering at the country level.

	(1)	(2)
	Civil war	Civil war
UDel temperature	0.042** (0.020)	-0.012 (0.019)
UDel temperature lag	0.004 (0.026)	-0.024 (0.041)
UDel precipitation	0.030 (0.028)	0.083 (0.079)
UDel precipitation lag	0.047* (0.026)	0.035 (0.034)
Constant	-0.997 (0.941)	1.084 (1.432)
Number of Observations	1056	288
R squared	0.645	0.592

Standard errors in parentheses
 * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Marshall Burke et al., 'Climate and Civil War: Is the Relationship Robust?' (National Bureau of Economic Research, 2010).

In response to this, Burke et al (2010) attribute this to increased economic growth and democratisation (as measured by the Polity Score) since 2000.

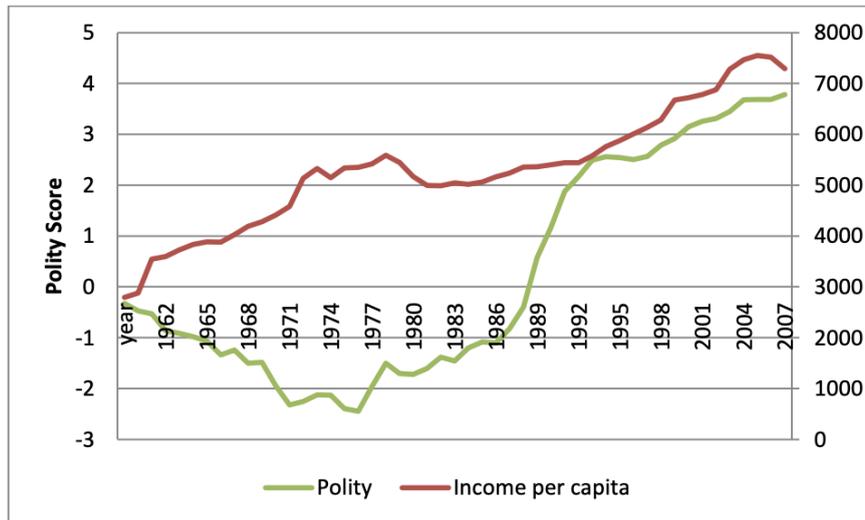
Table 5: Early and late period GDP per capita growth and Polity score, average across Sub-Saharan Africa. GDP data are from Penn World Tables (Heston et al 2009), and Polity data from the PolityIV Project (Marshall et al 2009)

	1981-2002	2003-2008
GDP per capita growth	0.5	2.8
Polity Score (-10 to 10)	-2.5	1.8

Source: Marshall Burke et al., 'Climate and Civil War: Is the Relationship Robust?' (National Bureau of Economic Research, 2010).

However, in Burke et al (2009), their optimistic scenario assumed that the effects of climate change would outweigh the effects of economic growth and democratisation, which is not what has happened. Their optimistic scenario assumed that the increase in democratisation followed trends from 1981-2002. But the increase in the polity score in Sub-Saharan Africa has actually been slower since 2002.

Figure 1: Global Democracy and Income



Source: Robert H. Bates, Ghada Fayad, and Anke Hoeffler, 'The State of Democracy in Sub-Saharan Africa', *International Area Studies Review* 15, no. 4 (2012): 323–38.

On the other hand, economic growth in reality was 0.8% above the assumed rate in their optimistic scenario. Still, this effect would not be large enough to produce the decline in conflict that is so far outside their model. Their bottom fifth percentile was a decline of 40% by 2030. In fact, conflict declined by 50% in only the next 6 years. Reality was very different to their predictive model.

Burke et al (2010) say:

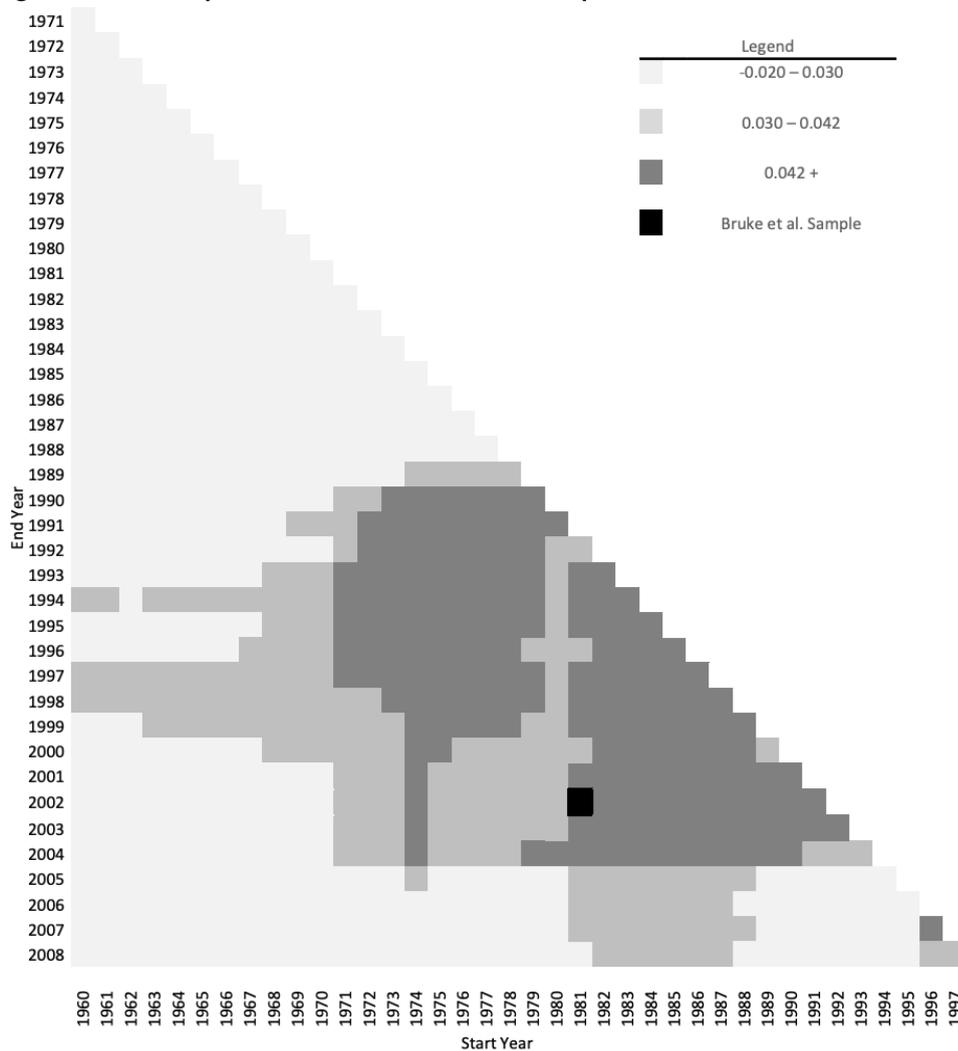
“Unlike our baseline climate dataset (CRU 2.1), the UDel data allow us to extend our analysis through 2008. We do not think that changing the study time period constitutes a “robustness test” but rather an answer to a slightly different question: i.e., were African economic and political institutions as sensitive to variation in climate over the most recent decade as they were over the previous two? Africa was clearly a different continent over the last decade compared to the two decades previous (Miguel 2008). As Table 5 shows, average annual per capita GDP growth over 2003-2008 was six times higher than the 1981-2002 period (where it was near zero), and the Polity Score (a -10 to +10 measure of democratic political institutions) improved an average of more than 4 points between the two periods.”⁵³⁹

Extending the dataset through to 2008 is not “answering a different question”: they themselves had asked this question in their original paper.

In a working paper, Buhaug et al illustrate the sensitivity of the Burke et al (2009) result to the time period studied. Burke et al (2009) study the effect of the period from 1981 to 2002. Here is how the coefficient would be affected by different date ranges:

⁵³⁹ Burke et al., 'Climate and Civil War'.

Figure 2. Size of temperature coefficient for various time periods since 1960



Source: Halvard Buhaug, Haavard Hegre, and Haavard Strand, ‘Sensitivity Analysis of Climate Variability and Civil War’, PRIO Paper, 2010.

A size of coefficient comparable to the Burke et al (2009) paper is found only in a certain cluster of date ranges, and is almost always (much) lower if the end date is after 2004. Longer time periods also tend to produce lower coefficients.

Sensitivity to civil war cut-off

As Buhaug notes, the finding in Burke et al (2009) is not one about civil war *onset*, but rather one about civil war *severity*, defined in a particular way. Burke et al (2009) measure the prevalence of civil war and count only conflict years that caused a minimum of 1,000 direct casualties. This approach has some strange implications:

“For example, consider the civil war in Sierra Leone. This conflict is widely accepted as lasting from March 1991 until the ceasefire and resulting Abuja Agreement in late 2000.* However, the Burke et al. article considers Sierra Leone at war in 1998–1999 only, the only 2 y in which direct annual casualty estimates crossed the 1,000 deaths threshold. Using climate statistics for 1997–1998 to explain a war that had caused

somewhere between 2,000 and 5,500 battle deaths by 1998 (14), however, makes little sense.”⁵⁴⁰

Other definitions of conflict seem reasonable, such as onset of civil wars of >1000 deaths, prevalence of >25 death conflicts, or battle deaths. If one uses any of these definitions, the result found in Burke is no longer statistically significant, and the mean effect is negative or close to zero.

Table 2. Testing for disagreement between results when alternative conflict variables are used

	Burke et al. (1) war years 1000+ (standardized)	Buhaug (8) model 5 incidence 1000+ (standardized)	Buhaug model 6 outbreak 1000+ (standardized)	Buhaug model 7 incidence 25+ (standardized)	Buhaug model 8 outbreak 25+ (standardized)	Buhaug model 9 outbreak 100+ (standardized)
Probability of occurrence	0.110	0.190	0.012	0.254	0.052	0.030
Temperature _t	0.390 (0.197)	-0.030 (0.110)	-0.408 (1.046)	0.060 (0.156)	-0.165 (0.504)	0.532 (0.790)
Temperature _{t-1}	0.120 (0.211)	-0.130 (0.147)	-0.755 (1.233)	-0.121 (0.128)	-0.083 (0.505)	-0.598 (0.581)
Precipitation _t	-0.209 (0.471)	0.326 (0.318)	-1.001 (4.212)	0.508 (0.281)	1.065 (1.316)	-0.455 (2.465)
Precipitation _{t-1}	0.227 (0.443)	0.296 (0.324)	0.205 (2.847)	0.093 (0.271)	0.352 (1.370)	-0.321 (2.017)
Observations	889	889	889	889	889	769
R-squared	0.657	0.765	0.090	0.652	0.130	0.099

Source: Solomon M. Hsiang and Kyle C. Meng, ‘Reconciling Disagreement over Climate–Conflict Results in Africa’, *Proceedings of the National Academy of Sciences* 111, no. 6 (11 February 2014): 2100–2103, <https://doi.org/10.1073/pnas.1316006111>.

(Note that these numbers are different to those found in Buhaug (2010) because that paper did not adjust for baseline risk of the different types of conflict).

Hsiang and Meng (2014) argue that these results are not necessarily inconsistent with those of Burke because the confidence intervals are wide enough to include the Burke results, so any difference could be due to sampling variation.⁵⁴¹ On the basis of this, they argue that Buhaug’s work is consistent with Burke’s results.

I don’t find this test plausible. Firstly, we should be interested in the mean effect, not whether an effect is within the 95% confidence interval of another effect. While we may not be able to rule out any difference with 95% confidence, the key question is whether, in expectation, this calls into question the posited causal link between climate change and conflict, which it does.

Secondly, it seems like what we should be interested in from a frequentist point of view is whether we can reject the null with 95% confidence, not whether we can reject a previous finding in the literature with 95% confidence. Hsiang and Meng’s sensitivity analysis shows

⁵⁴⁰ Buhaug, ‘Climate Not to Blame for African Civil Wars’.

⁵⁴¹ “In Table 2, we test whether the results in each of these models is different from the main result presented in Burke et al. (1) using seemingly unrelated regression (SUR), an approach that allows us to formally test whether or not two different regression models return statistically different results (14) (also 15, p. 153). Intuitively, this approach asks whether the “regression lines” describing the relationship between temperature and conflict in Burke et al.’s and Buhaug’s (8) analyses are statistically different from one another, while taking into account the fact that the studies are using related conflict outcomes where disturbances may be correlated” Solomon M. Hsiang and Kyle C. Meng, ‘Reconciling Disagreement over Climate–Conflict Results in Africa’, *Proceedings of the National Academy of Sciences* 111, no. 6 (11 February 2014): 2100–2103, <https://doi.org/10.1073/pnas.1316006111>.

that we can reject the null on other plausible operationalisations of conflict. Hsiang and Meng's approach makes the question of whether a finding has been refuted or not depend on the order in which papers are published. If Buhaug had published his paper first and found no effect, then Burke would have been unable to refute that with their data.

This shows that the Burke et al finding is highly sensitive to different definitions of conflict. Indeed, it seems like the measure we should care about is battle deaths in civil conflict. The relationship between climate and battle deaths is also noisy and uncertain. Contemporaneous temperature and lagged temperature each have a negative effect on battle deaths that is not statistically significant.

Table 1. Alternative model specifications

Specification	Model 1: War years 1,000+	Model 2: War years 500+	Model 3: War years 2,000+	Model 4: Battle deaths	Model 5: Log battle deaths
Temperature	0.044* (0.024)	0.008 (0.024)	0.003 (0.017)	-248.4 (261.4)	0.113 (0.222)
Temperature _{t-1}	0.010 (0.031)	-0.001 (0.035)	-0.008 (0.023)	-19.5 (268.7)	-0.120 (0.218)
Precipitation	-0.010 (0.070)	0.048 (0.072)	-0.042 (0.057)	-380.5 (690.6)	0.692 (0.503)
Precipitation _{t-1}	0.054 (0.051)	0.057 (0.075)	-0.052 (0.054)	-96.8 (711.6)	0.191 (0.506)
Intercept	-1.619 (1.214)	-0.511 (1.445)	0.233 (0.777)	7,350.4 (13,695.7)	-3.066 (10.118)
Country fixed effects	Yes	Yes	Yes	Yes	Yes
Country time trends	Yes	Yes	Yes	Yes	Yes
R ²	0.67	0.69	0.62	0.54	0.72
Civil war observations	125	173	85	226	226
Observations	889	889	889	889	889

Coefficients are ordinary least-squares regression estimates with country-fixed effects and time trends; SEs are in parentheses. All dependent variables generated from the same source (4). Model 1 is a replication of Model 2 of Burke et al. (1); slight differences in estimates are a result of different sources for the conflict variable. Models 2 and 3 apply alternative minimum casualty thresholds for civil war (500 and 2,000 annual deaths, respectively). Models 4 and 5 estimate the effect of climate on the annual number of battle deaths (count and logged estimates, respectively).

*P < 0.1.

Source: Buhaug, 'Reply to Burke et al.'

(Note that I'm not sure whether the coefficients in Models 2 and 3 are comparable to those in Model 1 because I'm not sure if they adjust for baseline risk. Nevertheless, again this does illustrate that if the thresholds of 500+ deaths or 2000+ deaths had been used, the effect of temperature on conflict would be ~0. This also raises some theoretical concerns about the posited mechanism by which climate change increases conflict. Burke et al (2009) argue that temperature causes conflict by affecting economic performance. It is not clear why climate damage to economic performance would cause conflicts with more than >1000 deaths, but not with >500 deaths or >2000 deaths.)

The coefficient for the number of battle deaths is negative (temperature reduces battle deaths) but the coefficient for log(battle deaths) is positive. This could happen because the number of deaths coefficient could be affected by extreme outliers, and the effect of outliers is limited with log transformation. Neither effect is statistically significant and there is a lot of noise.

Burke et al (2009) note that average battle deaths in conflicts in which >1000 died from 1981–2002 are 39,455 deaths/year.⁵⁴² However, average deaths per civil conflict in sub-Saharan Africa (so excluding North Africa) declined by a factor of 4 between the 1980s

⁵⁴² Burke et al., 'Warming Increases the Risk of Civil War in Africa', SI p3.

and the 2000s, despite warming.⁵⁴³ The chart below shows the trend in the number of civil wars and of battle deaths. There was a large decline after 1999 and since then the trend has been fairly stable.

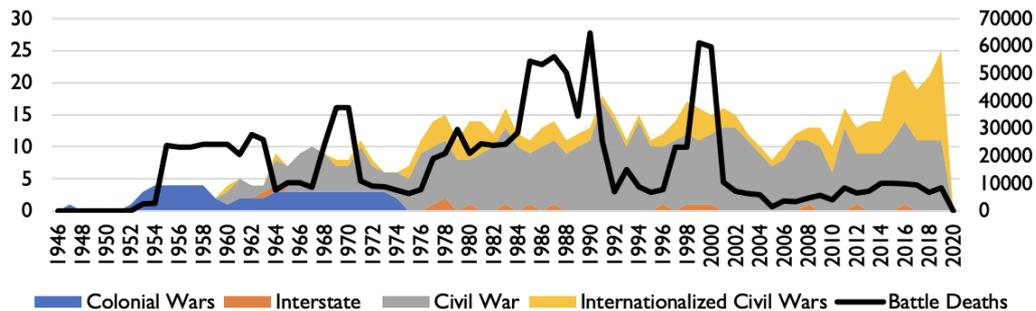


Figure 1: Battle deaths and state-based armed conflicts in Africa, by conflict type, 1946–2019

Source: UCDP/PRIO Armed Conflict Database, UCDP Battle Death Database (Pettersen & Öberg 2020) and Lacina and Gleditsch Battle Death Database

Source: Palik et al, ‘[Conflict Trends in Africa, 1989–2019](#)’, PRIO paper, 2020

Aggregated spatial and temporal scales

O’Loughlin et al (2014) write:

“The sub-Saharan Africa data used by Burke et al. (1) and the replications are countrywide data and averaged by year. Such a coarse spatial and temporal resolution limits our ability to uncover explanations for any relationships that emerge in statistical analyses. Increasingly in the field of conflict studies, researchers use subnational geographic data for single-country studies and cross national inquiries. Relying on a fine spatial resolution for the analysis of political violence allows intergroup social dynamics within a country to emerge. In the original Burke et al. data, countries as large as Sudan or Democratic Republic of Congo are single units, with political, economic, and climate characteristics that are uniform across enormous territories. This is a bold and naïve assumption, a view increasingly rejected by most conflict researchers because civil war tends to be concentrated in certain regions (e.g., southern Sudan and eastern Democratic Republic of Congo) (4).⁵⁴⁴

Qualitative analysis

The Burke et al (2009) result is highly dependent on six conflicts in Guinea-Bissau, Sierra Leone, Chad (in 1990), Congo, and Chad (in 1987). If these conflicts are removed, the coefficient falls close to zero.⁵⁴⁵ However, according to Buhaug et al, all of these conflicts involved foreign intervention. It might be that temperature influenced conflict in the country,

⁵⁴³ Buhaug, ‘Reply to Burke et Al.’

⁵⁴⁴ John O’Loughlin, Andrew M. Linke, and Frank D. W. Witmer, ‘Modeling and Data Choices Sway Conclusions about Climate-Conflict Links’, *Proceedings of the National Academy of Sciences* 111, no. 6 (11 February 2014): 2054–55, <https://doi.org/10.1073/pnas.1323417111>.

⁵⁴⁵ Buhaug, Hegre, and Strand, ‘Sensitivity Analysis of Climate Variability and Civil War’, 5.

which led to foreign intervention, but I would want to see a careful argument to that effect before putting much weight on the Burke et al (2009) study.

Because the result is driven by so few data points, qualitative analysis of the posited climate-economy-conflict link seems important. According to Buhaug et al:

“Another apparent commonality between these cases is the scarcity of references to climate anomalies and loss of agricultural income in news reports and narratives of the conflicts.”⁵⁴⁶

Overall view

I have a sceptical prior when dealing with individual social science studies that produce eye-catching and controversial findings using noisy data. My concerns especially centre around statistical significance, reporting bias and researcher degrees of freedom or the ‘[garden of forking paths](#)’. I think these problems will be especially severe for something as noisy as the relationship between climate and conflict. A deeper analysis of the Burke et al (2009) paper confirms that these concerns apply in this case.

This is important because, as I discuss below, most of the posited effect of climate change on conflict is driven by conflicts in Africa.

12.3.6. A review of the Burke et al (2015b) meta-analysis

The Berkeley economists have produced two meta-analyses on climate and conflict. I assume that the more recent Burke et al (2015b) paper is more representative of their current view, so will focus on that in what follows.

Intergroup conflict includes a lot

The meta-analyses showing a link between climate and intergroup conflict include quite different things in the definition of conflict, including:

- Civil war
- Communal violence involving 25 or more deaths
- Riots
- Coup probability
- Leadership exit
- Gang homicide
- Institutional change

This is important to bear in mind as the import of urban riots is quite different to the import of civil war. Nonetheless, most of the temperature studies concern civil conflict, involving >25 battle deaths. Note also that the effect sizes refer to the effect on the relative risk of each type of conflict, and so these types of conflict are not all aggregated.

Study inclusion criteria

In their primary research and in their meta-analyses, Burke, Hsiang and others rule out:

⁵⁴⁶ Buhaug, Hegre, and Strand, ‘Sensitivity Analysis of Climate Variability and Civil War’.

- Cross sectional studies that compare rates of conflict across different locations.
 - This risks omitted variable bias and over-controlling, as discussed in Chapter 10 on economic costs.⁵⁴⁷
- Studies that attempt to control for confounders, such as average income, institutional quality, and so on.
 - This can also lead to overcontrolling because populations differ in unobserved ways that become artificially correlated with climate when the control variable is included.⁵⁴⁸

Instead, they favour panel studies that promise superior causal identification. I agree with Burke et al that cross-sectional regressions should be interpreted with extreme caution, and that properly controlling for confounders is challenging. I would gloss their position as follows. Burke et al want to be able to make causal judgements with high confidence *on the basis of the study alone*, i.e. without introducing theoretical considerations or evidence from other domains.

However, I think that, from a Bayesian point of view, we should sometimes update from studies that require us to make theoretical judgements or use evidence from other domains. Thus, discarding all non-panel studies is too drastic. For example, income per head is strongly negatively correlated with conflict risk. It is true that income is also correlated with lots of other things that might affect conflict, but we also have a good theoretical explanation of why income might have a causal effect on conflict, so it would not be rational to discard such evidence. If income per head were not correlated with conflict risk, that would be a large update.

⁵⁴⁷ “One approach to the above problem would be to assume that populations or societies inhabiting different locations are identical to one another in all respects except their climate, usually after regression adjustment for observable economic, social, and political correlates of conflict. For example, Buhaug (2010a) compares the rate of civil war across different countries in Africa. Yet it seems implausible that the conditions needed for causal inference are met in this setting: There are many ways in which populations and societies differ from one another (e.g., culture, history), many of them unobserved or hard to measure, so we cannot infer whether a climatic treatment has a causal effect (Wooldridge 2002, Angrist & Pischke 2008). In the above example, the cross-sectional analysis by Buhaug (2010a) compares average rates of civil conflict in South Africa and Nigeria (among many comparisons), attributing observed differences to the different climates of these countries—despite the many other potentially important ways in which these countries differ.” Marshall Burke, Solomon M. Hsiang, and Edward Miguel, ‘Climate and Conflict’, *Annual Review of Economics* 7, no. 1 (2015): 577–617, <https://doi.org/10.1146/annurev-economics-080614-115430>.

⁵⁴⁸ “Some studies expand Equations 3 and 4 to explicitly control for potential confounding factors, such as average national income. For example, Buhaug (2010a) alters the analysis of a temperature–war association studied by Burke et al. (2009) to include indices for political exclusion and average income. Although well intentioned, this approach may introduce bias in the coefficients describing the effect of climate on conflict because these controls are endogenously determined and may themselves be affected by climate variation. This can cause the signal in the climate variable of interest to be inappropriately absorbed by the control variable or the estimate to be biased because societies differ in unobserved ways that become artificially correlated with climate when the control variable is included. This approach is commonly termed bad control (Angrist & Pischke 2008) and is a particular difficulty in this setting because climatic variables may affect so many of the socioeconomic factors commonly included as control variables, such as crop production, infant mortality, population (via migration or mortality), and even political regime type. To the extent that these outcome variables are used as covariates, studies might draw mistaken conclusions about the relationship between climate and conflict... In what follows, we modify estimates that rely on this method by excluding bad controls in our reanalysis.” Burke, Hsiang, and Miguel.

Moreover, showing that a relationship *could be* confounded does not show that it is *likely to be* confounded or that the effect is large. For example, it is true that climate change might influence average income and so if we control for average income, the signal from the climate variable would be inappropriately absorbed by the control variable. However, we might also be able to establish with other types of evidence that climate change is at best a minor influence on average income. In that case, controlling for average income would still reduce the bias of the estimate.

Study inclusion and coding decisions

I have doubts about some of the study inclusion and coding decisions in Burke et al (2015b).

Firstly, this may not be particularly important, but it is somewhat unclear how to disaggregate the findings in different papers included in the meta-analysis. Burke et al (2015b) include Burke et al (2009) discussed above, as well as the Buhaug (2010) finding, which I think is the one using the >25 battle death definition of conflict. Burke et al (2015b) counts this as two contrasting findings, but there were many other findings in the back and forth regarding Burke et al (2009), including null effect at >500 death or >2000 death thresholds or for battle deaths. Following the logic of including one of the Buhaug (2010) results, one could include multiple different null results. This problem arises from the expansive definition of intergroup conflict. I'm not sure what the correct solution to this problem is, but it is worth noting.

Secondly, Burke et al (2015b) only includes the effect found in Burke et al (2009) on temperature, but not on precipitation. But Burke et al (2010) notes that extending the years of conflict to 2008, precipitation has almost zero effect on conflict prevalence, and that the effect is not statistically significant.⁵⁴⁹ The null finding for precipitation does not make it into the meta-analysis, even though Burke et al (2009), on which Burke et al (2010) is based, is included in the meta-analysis.

Finally, I have looked into some of the other included studies and I find some of the coding decisions difficult to understand. For example, Brückner and Ciccone (2011) finds that “negative rainfall shocks lead to *significant democratic improvement and, in particular, a tightening of executive constraints, greater political competition, and more open and competitive executive recruitment*” [my emphasis].⁵⁵⁰ Burke et al (2015b) code this as a form of intergroup conflict - ‘institutional change’ specifically. The terms ‘death’, ‘fatalities’, ‘casualties’, ‘oppression’ and their cognates are not used in the Brückner and Ciccone paper, and ‘conflict’ is only mentioned in discussion of the methodology of other papers in a passing footnote.⁵⁵¹ The paper is just not about intergroup conflict, and in fact shows that weather variability has positive effects.

A 2010 study by Paul Burke (different to Marshall Burke, who wrote the meta-analysis) and Andrew Leigh finds that weather shocks to GDP *increase* the chance of democratisation.⁵⁵²

⁵⁴⁹ Burke et al., ‘Climate and Civil War’, Table 1.

⁵⁵⁰ Markus Brückner and Antonio Ciccone, ‘Rain and the Democratic Window of Opportunity’, *Econometrica* 79, no. 3 (2011): 923–47.

⁵⁵¹ Brückner and Ciccone, fn4.

⁵⁵² Paul J. Burke and Andrew Leigh, ‘Do Output Contractions Trigger Democratic Change?’, *American Economic Journal: Macroeconomics* 2, no. 4 (2010): 124–57.

Burke et al (2015b) again code this as a form of intergroup conflict involving institutional change. Burke and Leigh (2010) note that

“The result that stronger economic growth reduces the short-run likelihood of democratic change needs to be considered alongside the finding of Miguel, Satyanath, and Sergenti (2004) that stronger growth reduces the short-run likelihood of civil conflict in Africa. The combined evidence indicates that output contractions trigger anti-government resistance activities. Further research into the conditions which determine whether output contractions lead to peaceful or to violent opposition to incumbent regimes may be of considerable value.”

Burke and Leigh (2010) is not about violence or conflict and in fact shows that weather shocks have positive social effects.

A 2012 study by Paul Burke on the effects of weather variability on leadership exits is included. There are two findings in the paper: one for *regular exits* which occur when a leader leaves office according to the prevailing rules, provisions, conventions and norms of the country; and one for *irregular exits*, which occur when a leader is removed from office in contravention of rules and conventions (for example, by coups, assassinations, military power struggles, or removal by domestic rebel forces or revolts). Only the latter kind of exit can be classed as demonstrating intergroup conflict. However, as I understand it, though I am not sure, the Burke et al (2015b) meta-analysis includes the effect of the weather on both regular and irregular forms of political exit.⁵⁵³

These are all important errors and they seem to point towards a bias in favour of including studies that find a positive effect on conflict and excluding studies that find a negative effect.

Reporting bias

Burke et al (2015b) demonstrate evidence of reporting bias in the papers included in their meta-analysis, but they do not adjust for this when making their estimates of the effect of conflict. They use the following tests of reporting bias:

- A.** If there is no reporting bias, across the set of studies the statistical power of studies should increase with their sample size.
- B.** If there is large reporting bias, then large sample sizes should provide no benefit in terms of statistical power. (Maximal reporting bias would be when sample size and statistical power are perfectly anti-correlated).

Eventuality A is shown by the 45 degree diagonal line on the figure below, and eventuality B is shown on the horizontal line. (Maximal reporting bias would be shown by a downward sloping 45 degree line).

⁵⁵³ Compare Figure 3 of Burke et al (2015) with Table 12 of Paul Burke (2012). The latter shows a small and non-significant effect for irregular exit, whereas Burke et al (2015) shows a larger marginally significant effect.

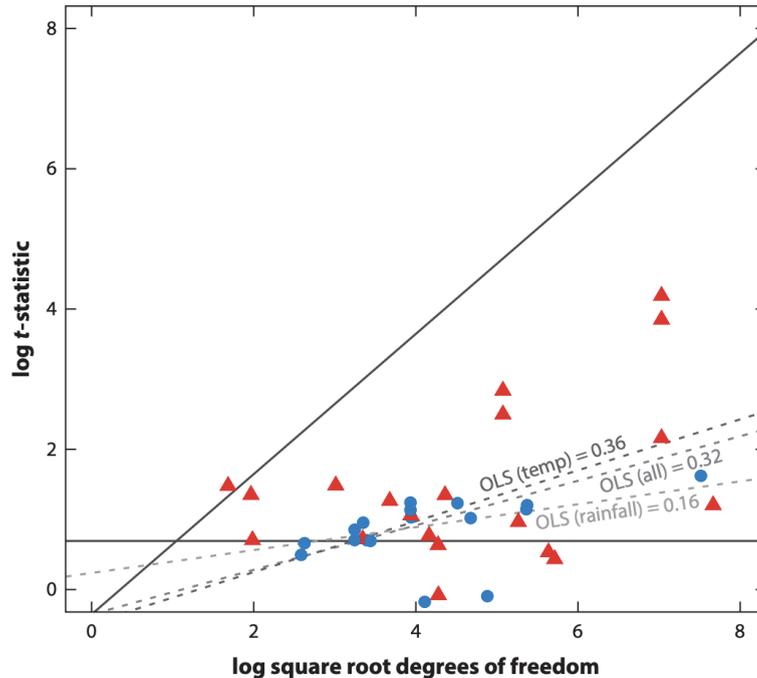


Figure 5

Relationship between the log of the t -statistic and the log of the square root of the degrees of freedom, using author-reported t -statistics. Circles represent studies focusing on rainfall, and triangles represent studies focusing on temperature. Ordinary least squares (OLS) estimates of the slope of the relationships for all estimates, temperature-only estimates, and rainfall-only estimates are given by the dashed lines. The solid 45° line (the unit elasticity) indicates the slope of 1 that theory predicts would occur in the absence of any publication bias.

Source: Burke, Hsiang, and Miguel, 'Climate and Conflict'.

As this shows, the trend line across studies for temperature is much closer to the horizontal than the 45 degree line: the correlation between sample size and statistical power is quite weak (0.36 in the case of temperature). Burke et al (2015b) say this about this result:

“We strongly reject a slope of zero for both the full sample (P value < 0.01) and the temperature subsample (P value < 0.05) and marginally reject a zero slope for the precipitation estimates (P value < 0.10). And although we can also reject a one-to-one relationship for each sample, studies with larger sample sizes on average do have larger t -statistics in the climate and conflict literature we survey, suggesting that authors with large samples are not simply searching through specifications or data mining to find significant effects at exactly the 95% confidence interval. We note that for both samples, the upward relationship stands in sharp contrast to the results of Card & Krueger (1995), with the negative slope they estimate. Our estimates are more similar to that of Disdier & Head (2008), who interpret their results as ruling out any large role for publication bias in the trade literature they survey.”

Firstly, just as they strongly reject a slope of zero, they even more strongly reject a slope of 45 degrees. Contra Burke et al, this is indeed evidence that authors with large samples are searching through specifications or data mining to find significant effects at the 95% confidence interval. Finding that the correlation between sample size and statistical power is not zero is a low bar test for reporting bias.

Secondly, I don't think the relevant benchmark for reporting bias adjustment should be the practice of other economists.

Given that they found clear reporting bias, Burke et al should have adjusted for it in their meta-analysis.

Implausibly precise estimates

The meta-analytic finding in Burke et al (2015b) is as follows:

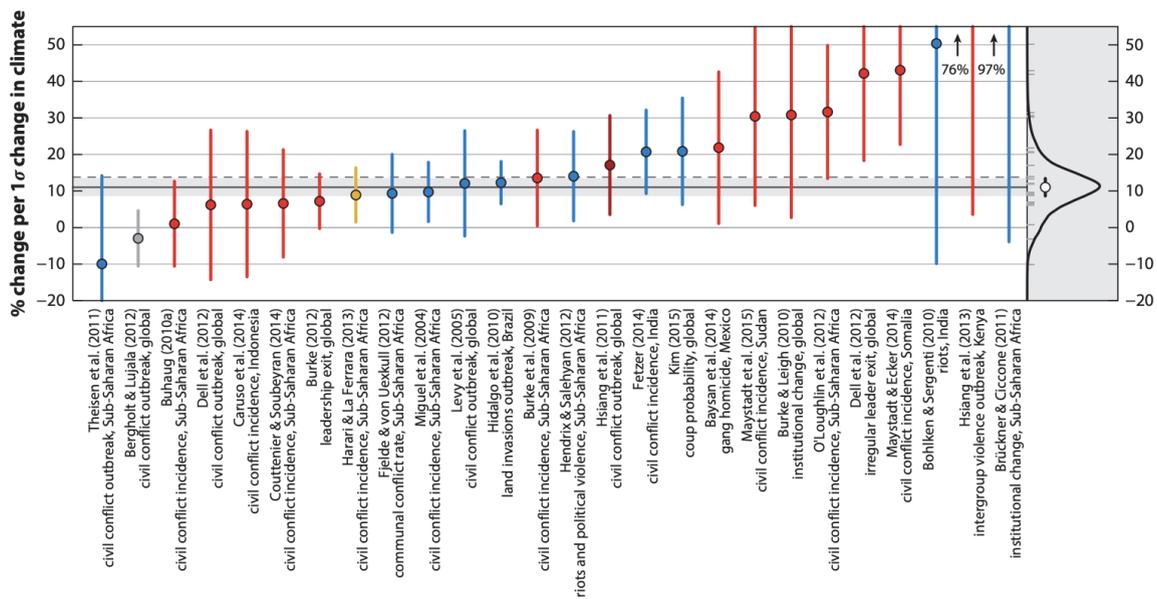


Figure 3

Estimates for the effect of climatic events on the risk of intergroup conflict using authors' preferred specifications. Each marker represents the estimated effect of a 1σ increase in a climate variable, expressed as a percentage change in the conflict variable relative to its mean, and whiskers represent the 95% confidence interval. Colors indicate temperature (red), rainfall loss (blue), drought (yellow), the El Niño Southern Oscillation (dark red), and storms (gray). The dashed line is the median estimate, and the solid dark gray line is the precision-weighted mean with its 95% confidence interval shown in gray. The panel on the right shows the precision-weighted mean effect (circle) and the distribution of study results (gray ticks); probability distributions are the posterior for the expected distribution of an additional study (solid black line).

The 95% confidence interval for the mean effect is shown in the grey box on the right by the white circle with whiskers. This means that Burke et al (2015b) are extremely confident that the true effect per standard deviation increase in climate change is between 9 and 13%; it is almost impossible to be <7% or >16%. This seems implausibly precise for two things as noisy, variable, multi-causal and complex as climate change and conflict.

Indeed, they find an effect of precipitation on conflict that is (a) tiny - a 3.5% increase in relative risk per standard deviation change; and (b) statistically significant. This result is difficult to believe.

I think what is going on here is that they are estimating an effect size on the assumption that their model is correct, but in reality there is much more uncertainty about their model than there is uncertainty conditional on their model.

One can visually see the effect of publication bias in the chart above. The meta-analytic effect size is meant to be 3.5% for precipitation, so we should expect to see precipitation

studies dotted either side of zero. Instead, the included precipitation studies (in blue) overwhelmingly find a positive effect, with numerous studies finding an effect of more than 10%. The mean estimated effect of temperature (in red) on conflict is 11%, but there are no studies finding a negative effect and most of the studies find an effect greater than 20%. One would expect to see at least some negative studies for an effect this small in comparison to other factors, solely due to noise. (For comparison, civil conflicts in Africa declined by 50% between 2003 and 2008, and deaths per conflict declined by 400% in Sub-Saharan Africa between the 1980s and the 2000s).

Africa focus

In Burke et al (2015), if we include studies on the El Nino Southern Oscillation, there are 10 studies of the effect of temperature on civil conflict, 7 of which are in Africa, 2 global and 1 focuses on Indonesia. I now discuss the non-Africa focused studies.

Caruso et al (2016) - Climate Change, Weather Shocks and Violent Conflict: A Critical Look at the Evidence

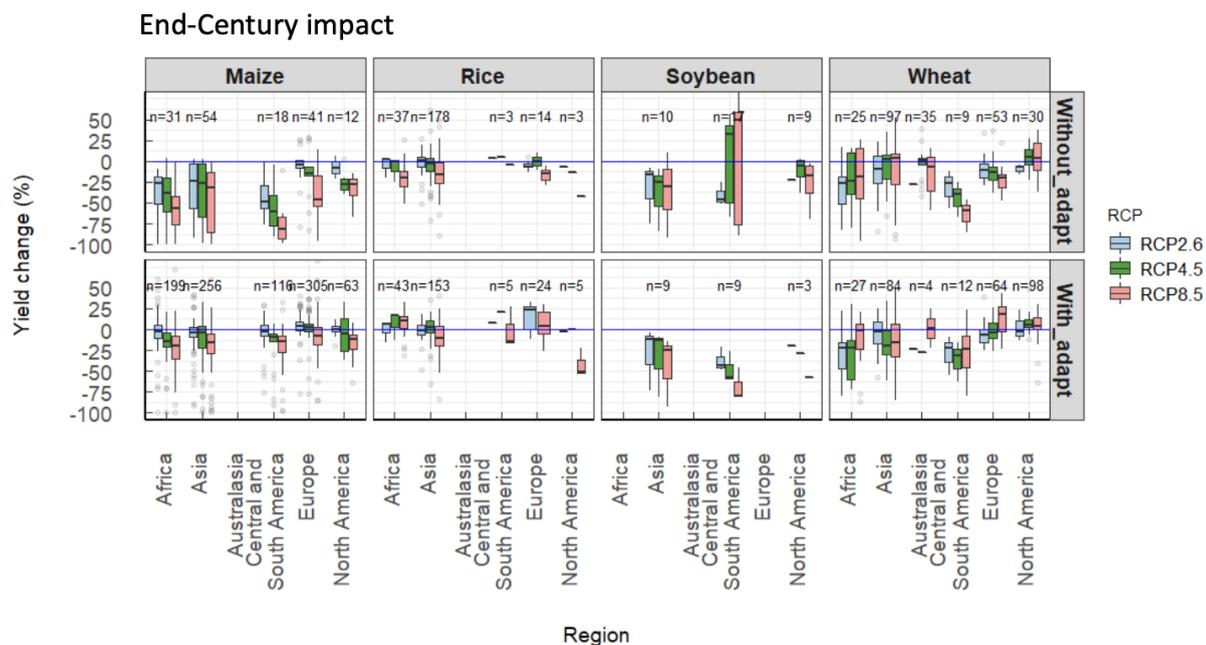
Caruso et al (2016) studies the effect of climate change on rice production and subsequent violence in Indonesia over a ten year period from 1993 to 2003.⁵⁵⁴ Although Burke et al code this study as being about civil conflict, which is standardly defined as conflict between the state and others involving >25 battle deaths, it is actually about something more expansive:

“violence perpetrated by a group on another group (as in riots), by a group on an individual (as in lynchings), by an individual on a group (as in terrorist acts), by the state on a group, or by a group on organs or agencies of the state.”⁵⁵⁵

I don't put much weight on this study. Firstly, the study period is only ten years, so it is hard to know whether the result is robust. Secondly, climate change is projected to have mild effects on rice yield, with adaptation, on average, in both temperate and tropical regions.

⁵⁵⁴ Burke et al (2015) discusses the earlier working paper of the Caruso et al (2016) paper.

⁵⁵⁵ Raul Caruso, Ilaria Petrarca, and Roberto Ricciuti, 'Climate Change, Rice Crops, and Violence: Evidence from Indonesia', *Journal of Peace Research* 53, no. 1 (1 January 2016): 66–83, <https://doi.org/10.1177/0022343315616061>.



Supplementary Fig. S3 Climate change impacts on four crops in the mid 21st century with and without adaptation in IPCC regions by regions at mid-century (MC, 2040-2069, upper panels) and end-century (EC, 2070-2100, lower panels). n is the number of simulations. The box is the interquartile range (IQR) and the middle line in the box represents the median. The upper- and lower-end of whiskers are median $1.5 \times \text{IQR} \pm \text{median}$. Open circles are values outside the $1.5 \times \text{IQR}$.

Source: Toshihiro Hasegawa et al., 'A Global Dataset for the Projected Impacts of Climate Change on Four Major Crops', *Scientific Data* 9, no. 1 (16 February 2022)
<https://doi.org/10.1038/s41597-022-01150-7>.

This suggests that the mechanism of climate change => damaged rice production in Asia => civil conflict is extremely weak. For comparison, rice yields have increased by more than 200% over the last 60 years.

Dell et al (2012) - Temperature Shocks and Economic Growth: Evidence from the Last Half Century', *American Economic Journal: Macroeconomics*

The findings of the only global temperature study on civil conflict - Dell et al (2012) - are shown below (columns 5 and 6):

TABLE 6—POLITICAL ECONOMY EFFECTS

	Any change in POLITY score (1)	Leader transition (2)	Regular leader transition (3)	Irregular leader transition (4)	Start of conflicts (conditional on conflict = 0 in $t - 1$) (5)	End of conflicts (conditional on conflict > 0 in $t - 1$) (6)
Temperature	-0.013 (0.009)	-0.002 (0.015)	0.004 (0.015)	-0.005 (0.004)	-0.006 (0.006)	0.005 (0.060)
Temperature × Poor	0.040** (0.016)	0.033 (0.023)	-0.017 (0.017)	0.050*** (0.013)	0.012 (0.013)	0.003 (0.068)
Precipitation	0.001 (0.003)	0.003 (0.002)	0.003 (0.003)	0.000 (0.001)	0.000 (0.001)	0.023 (0.019)
Precipitation × Poor	0.008 (0.005)	-0.008* (0.004)	-0.008** (0.004)	0.000 (0.002)	0.000 (0.002)	-0.031 (0.020)
Observations	5,388	6,624	6,624	6,624	5,702	852
R^2	0.14	0.18	0.2	0.11	0.09	0.43
Within R^2	0.003	0.001	0.001	0.004	0.000	0.004
Temperature effect in poor countries	0.027* (0.015)	0.031* (0.017)	-0.013 (0.010)	0.044*** (0.013)	0.007 (0.011)	0.008 (0.037)
Precipitation effect in poor countries	-0.009** (0.004)	-0.005 (0.004)	-0.005* (0.003)	0.000 (0.002)	0.000 (0.002)	-0.009 (0.007)

Notes: Column 1 uses data from the POLITY IV dataset; columns 2–4 use data from the Archigos dataset; and columns 5–6 use data from the PRIO dataset. Columns 1–4 include country FE, region × year FE, and poor × year FE; columns 5–6 include country FE and year FE. Robust standard errors are in parentheses, adjusted for two-way clustering at parent-country and year-region levels. Sample includes all countries with at least 20 years of WDI growth observations (i.e., the same set of countries considered in the previous tables).

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Source: Dell, Jones, and Olken, 'Temperature Shocks and Economic Growth'.

This shows that at the global level, the mean effect of temperature is to reduce civil conflict outbreak (though the effect is not statistically significant), but in poor countries (with average income below the median - therefore mainly countries in Africa and South Asia), temperature increases civil conflict (though the effect is not statistically significant).

Hsiang et al (2011) - Civil conflicts are associated with the global climate, Nature

Hsiang et al (2011) is a global study that argues that the warmer and dryer conditions associated with El Nino events are associated with civil conflict onset. However, Klomp and Bulte (2013) do not find an El Nino effect. Klomp and Bulte (2013) is not included in the Burke et al (2015b) meta-analysis, so perhaps Klomp and Bulte use an excluded methodological approach. However, they do say that a "When we further probe the robustness of the Hsiang et al. result and use their exact model specification and statistical program in a rolling regression framework (so not the standard logit model), we also do not find any robust result of the ENSO effect in Africa or Asia".⁵⁵⁶ I'm not sure what explains that.

All of this suggests that it would be premature to infer anything from Burke et al (2015b) about the risk of intergroup conflict from increased temperature *outside of Africa*.

⁵⁵⁶ Jeroen Klomp and Erwin Bulte, 'Climate Change, Weather Shocks, and Violent Conflict: A Critical Look at the Evidence', *Agricultural Economics* 44, no. s1 (2013): 63–78.

This is one instance of a more general problem in the literature, which is that research overwhelmingly focuses on Africa, which is not necessarily where climate stress is greatest. Adams et al note:

“In contrast, the sampling of countries to be studied seems to be barely informed by the independent variable. A high exposure and a high vulnerability to climate change according to the ND-GAIN index²³ are negatively, but not significantly, correlated with the number of times a country is mentioned.”⁵⁵⁷

Table 3 | Countries most often mentioned in the climate–conflict literature compared with the countries most exposed to and at risk from climate change

Rank	Number of mentions	Exposure score	Climate Risk Index
1	Kenya	11	Rwanda 0.622 Honduras 11.33
2	Sudan	11	Kiribati 0.620 Myanmar 14.17
3	Egypt	8	Burundi 0.617 Haiti 18.17
4	India	7	Zambia 0.613 Nicaragua 19.17
5	Nigeria	7	Tuvalu 0.612 Philippines 21.33
6	Syria	7	Marshall Islands 0.600 Bangladesh 25.00
7	Israel/Palestine	6	Yemen 0.597 Pakistan 30.50
8	Ethiopia	5	Seychelles 0.582 Vietnam 31.33
9	Iraq	5	Oman 0.568 Guatemala 33.83
10	South Sudan	5	Micronesia 0.567 Thailand 34.83

Hsiang and Burke [respond](#) to this

“Adams et al.’s error arises because they confuse sampling observations within a given study based on the dependent variable (a major statistical violation) with the observation that there are more studies in locations where the average of a dependent variable, the conflict rate, is higher (not a violation). Nowhere does Adams et al. provide evidence that any prior analysis contained actual statistical errors.”

The problem with this is that Burke and Hsiang make claims about a global increase in conflict risk due to climate change. Take this extreme example. Suppose all of the literature focused on conflict in Africa. Even if all of this literature was methodologically sound, it would be dubious to produce a meta-analysis saying that climate has global effects on conflict because you would have ignored many of the data points where (i) climate stress is also high and (ii) rates of conflict may be different.

The neglect of Asia is especially important because, as of 2014, Asia was home to nearly half of the world’s active civil wars, and had the highest density of armed conflicts per country in the world.⁵⁵⁸

⁵⁵⁷ Adams et al., ‘Sampling Bias in Climate–Conflict Research’. See also Cullen S. Hendrix, ‘The Streetlight Effect in Climate Change Research on Africa’, *Global Environmental Change* 43 (1 March 2017): 137–47, <https://doi.org/10.1016/j.gloenvcha.2017.01.009>.

⁵⁵⁸ Gerdis Wischnath and Halvard Buhaug, ‘On Climate Variability and Civil War in Asia’, *Climatic Change* 122, no. 4 (1 February 2014): 710, <https://doi.org/10.1007/s10584-013-1004-0>.

Other issues with panel studies

Panel studies on conflict have the same problem as the panel studies for economic costs, discussed in Chapter 10. Climate change is different to inter-annual weather variation. It is not clear in which direction this biases the finding.

However, panel data does allow good causal identification. Buhaug and others typically want to control for things like GDP and institutions, which introduces the risk of omitted variable bias and overcontrolling.

Overall judgement on these estimates

I think the Burke and Hsiang meta-analytic estimates of the effect of climate on conflict are substantial overestimates. In other domains, it has been shown that meta-analyses *that try to correct for publication bias* typically substantially overestimate effect sizes.

- In psychology, Kvarven et al (2020) find that meta-analytic effect sizes are nearly three times greater than those found in large pre-registered randomised control trials, even accounting for meta-analytic methods that try to adjust for publication bias.⁵⁵⁹
- Simulations suggest that meta-analyses that try to control for reporting bias using trim and fill will typically overestimate effect sizes, usually by a factor of 2 or more.⁵⁶⁰
- In medicine, where there are requirements for pre-registering large RCTs, there is a false positive and false negative rate of one third in meta-analyses.⁵⁶¹

Despite evidence of publication bias, Hsiang and Burke's meta-analyses do not try to correct for publication bias. Given information from other fields, my best guess is that, due to publication bias, their meta-analytic effect sizes are likely overestimated by a factor of more than 2, which suggests that the effect of climate change on civil conflict is <5% per standard deviation. The other factors I have discussed above also suggest that their estimates are too high.

This puts the adjusted estimates more in the ballpark of other quantitative estimates identified by Sakaguchi et al (2019) in their systematic review.⁵⁶²

⁵⁵⁹ Amanda Kvarven, Eirik Strømland, and Magnus Johannesson, 'Comparing Meta-Analyses and Preregistered Multiple-Laboratory Replication Projects', *Nature Human Behaviour*, 23 December 2019, 1–12, <https://doi.org/10.1038/s41562-019-0787-z>.

⁵⁶⁰ Uri Simonsohn, Leif D. Nelson, and Joseph P. Simmons, 'P-Curve and Effect Size: Correcting for Publication Bias Using Only Significant Results', *Perspectives on Psychological Science* 9, no. 6 (2014): fig. 3.

⁵⁶¹ "According to our analysis, if there had been no subsequent randomized, controlled trial, the metaanalysis would have led to the adoption of an ineffective treatment in 32 percent of cases (100 percent minus the positive predictive value) and to the rejection of a useful treatment in 33 percent of cases (100 percent minus the negative predictive value)." Jacques LeLorier et al., 'Discrepancies between Meta-Analyses and Subsequent Large Randomized, Controlled Trials', *New England Journal of Medicine* 337, no. 8 (21 August 1997): 536–42, <https://doi.org/10.1056/NEJM199708213370806>.

⁵⁶² "Of the forty-three studies that found a positive relationship, some studies report results in terms of a percentage change in conflict risk (Burke et al. 2009; Hsiang, Meng, and Cane 2011; Fjelde and von Uexkull 2012; Hendrix and Salehyan 2012). They find the risk of conflict rises between 0.5 percent and 4.5 percent relative to a unit change in the climate indicator." Sakaguchi, Varughese, and Auld, 'Climate Wars?'

Overall, I think that if the Burke et al (2015b) meta-analysis properly adjusted for study inclusion and publication bias, they would find that the panel studies would, like the studies using other methodologies, also find a mixture of results on the climate-conflict link.

Implications of the estimates

Burke et al estimates taken at face value

Suppose that you take the Burke et al (2015b) estimates at face value and say that the risk of intergroup conflict in Africa will increase by 11% for each SD increase in warming. Assume that by 2100, each region will experience a 8 standard deviation increase in temperature, which is what we might expect in a 5 degree world.

Since 2000, battle deaths in Africa have [been](#) on the order of 5,000 per year.

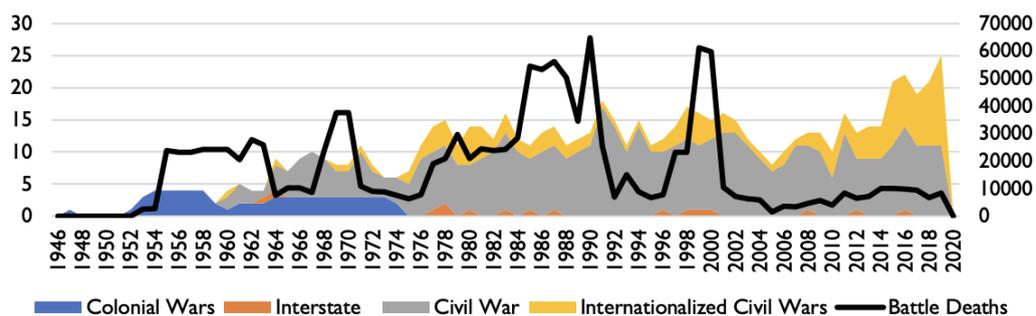


Figure 1: Battle deaths and state-based armed conflicts in Africa, by conflict type, 1946–2019

Source: UCDP/PRIO Armed Conflict Database, UCDP Battle Death Database (Pettersson & Öberg 2020) and Lacina and Gleditsch Battle Death Database

Source: Palik et al, '[Conflict Trends in Africa, 1989–2019](#)', PRIO paper, 2020

This suggests that battle deaths will increase to 40,000 by 2100, *other things equal*. This is bad, but much smaller than other public health problems. For instance, more than a million people die in road accidents every year.

In essence, this would return Africa to the civil conflict prevalence and severity seen in the early 1990s.

My favoured estimates

As I have said, I think the Burke and Hsiang meta-analyses overestimate the effect of climate change on conflict. If the effect on conflict is a 5% increase per standard deviation change, then we end up with civil conflict rising by 40% by 2100, which would increase deaths to around 6,000 by 2100, *other things equal*.

Mechanisms of impact

Burke et al (2015b) discuss several potential mechanisms by which climate could affect conflict, including:⁵⁶³

- Damaging economic output

⁵⁶³ Burke, Hsiang, and Miguel, 'Climate and Conflict', sec. 3.

- Affecting population density by increasing rural to urban migration
- Climate change might affect the chance of a successful attack
 - This might be due to logistics - for example floods might affect the road available to military vehicles
 - Or it might be because economic damage reduces the security of incumbent leaders, as illustrated by the evidence that climate change affects leadership exit.
- Temperature might increase aggression via a physiological mechanism.

Burke et al (2015b) believe that the psychological mechanism is the most plausible because economic conditions and most other plausible alternative mechanisms usually do not respond to the climatic conditions on the timescale analysed.

I do not find it plausible that climate change could increase conflict risk by systematically increasing the probability of a successful attack. Firstly, as discussed above, Burke et al (2015b) seem to misinterpret the evidence on leadership exit, which shows noisy effects close to zero. Indeed, many studies suggest that climate change increases democratisation, which is widely agreed to reduce conflict risk. Secondly, climate change is set to alter precipitation patterns, which will mean more floods in some locations and less in others. The effects of climate change on logistics also seem to be at most a small driver of conflict.

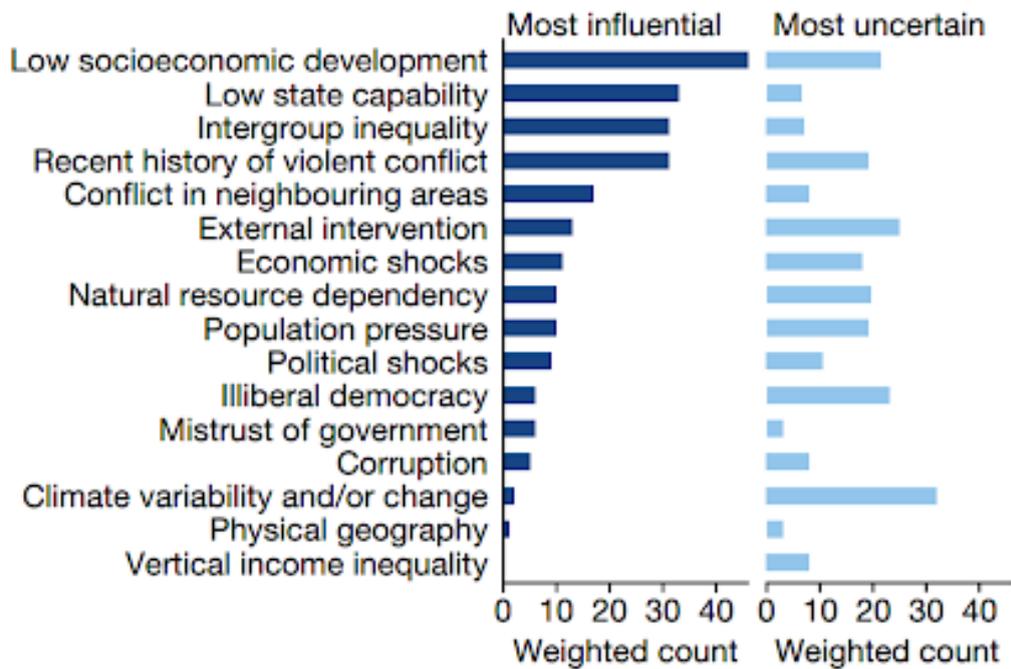
The effect on rural-urban migration raises some interesting questions. As I discuss below, population size is a correlate of civil conflict risk, but population growth is also correlated with reduced *per capita* conflict risk.

Thus, it seems that the economic mechanism and the direct physiological mechanism seem to be the most plausible candidates among those raised by Burke et al. I discuss the direct physiological mechanism in the next Chapter.

12.3.7. Climate less important than other factors

Despite disagreement in the literature, all major players in the debate seem to agree that climate change has been much less important than other factors as an influence on conflict.

An expert elicitation study including Buhaug and Burke found almost no-one putting climate variability as the most influential determinant of war so far, with socioeconomic development, state capability, intergroup inequality and a recent history of violent conflict being the most important.



Source: Katharine J. Mach et al., 'Climate as a Risk Factor for Armed Conflict', *Nature* 571, no. 7764 (July 2019): 193–97, <https://doi.org/10.1038/s41586-019-1300-6>.

Most experts held that around 10% of civil conflict risk has been related to climate change so far.

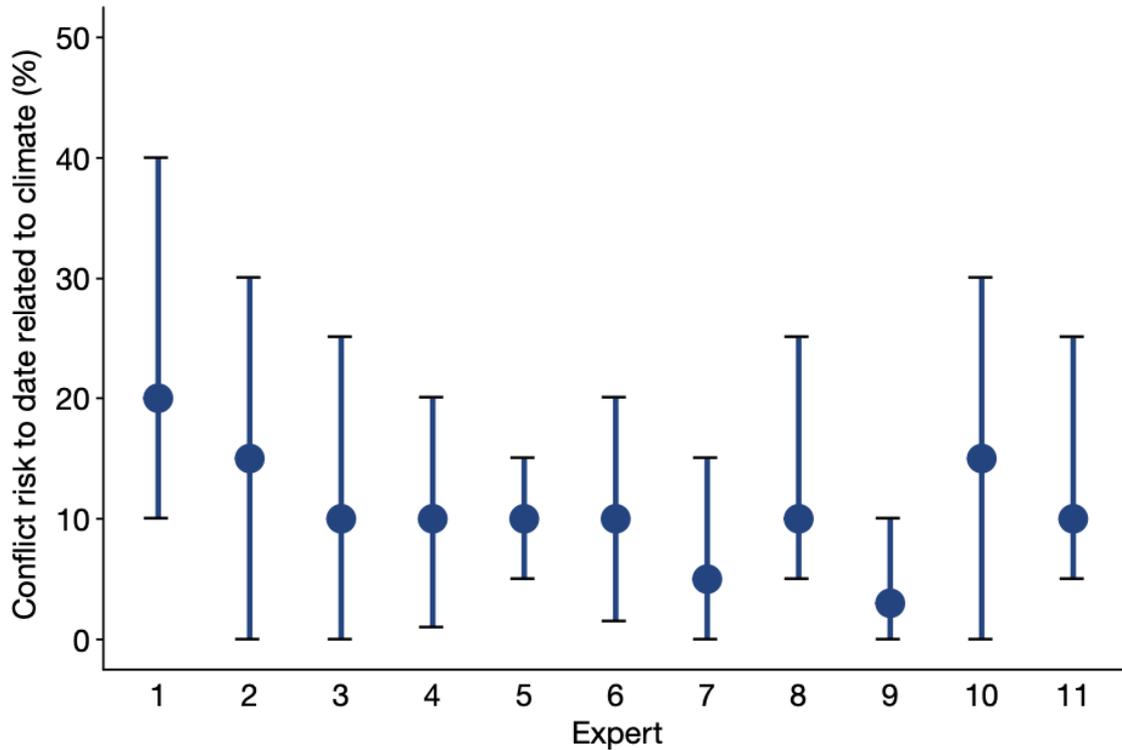


Fig. 1 | The estimated relationship between climate and conflict risk to date. Each expert provided subjective probabilistic judgments of the percentage of total conflict risk related to climate across experiences over the past century. The estimated 1st, 50th and 99th percentiles are shown for each expert.

Looking to the future, most of the surveyed experts felt that, assuming current levels of socioeconomic development, population and government capacity, 2°C would have a negligible or moderate effect on civil conflict, with a plurality finding negative change. For 4°C, there would be a moderate increase. A ‘substantial’ increase is defined in the elicitation as ‘involving severe and widespread effects’. Since it is not clear what the definition means, it is not clear how to interpret the chart below:

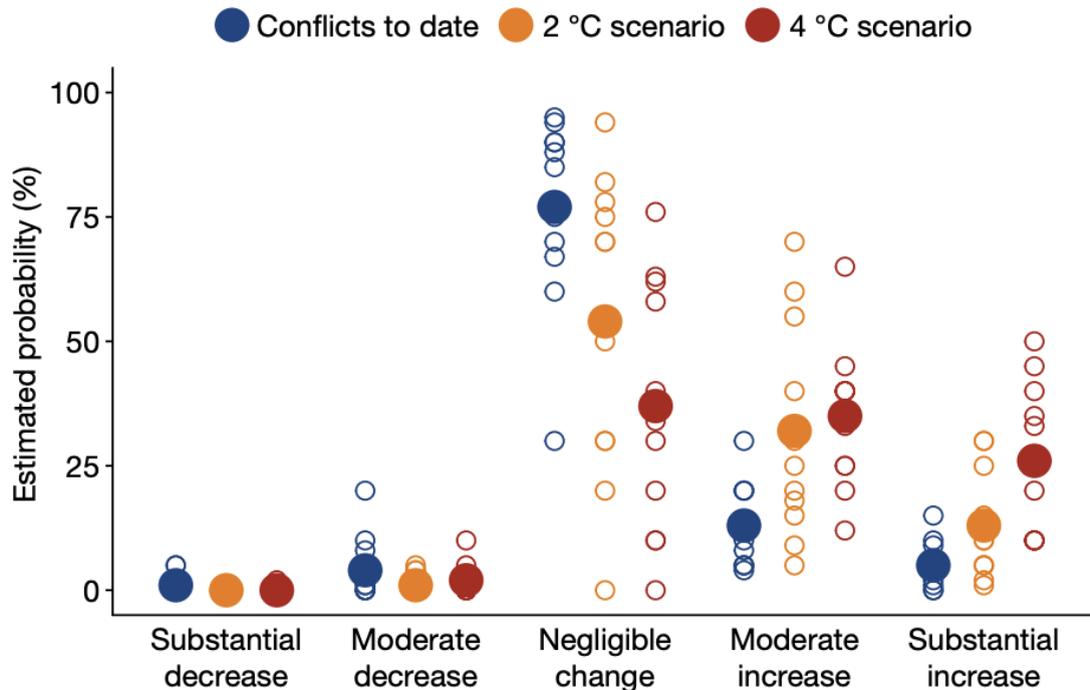


Fig. 2 | Estimated changes in the relationship between climate and conflict risk under increasing climate change. For three scenarios, each expert estimated the likelihood that climate leads to negligible, moderate or substantial changes in conflict risk. For violent conflicts to date (blue), probability estimates indicate how frequently climate variability and change have led to the specified changes in conflict risk. For the approximately 2 °C (orange) and approximately 4 °C (red) warming scenarios, probability estimates indicate potential changes in conflict risk compared to the current climate. For these hypothetical 2 °C and 4 °C scenarios, each expert considered associated effects of climate change for current societies, assuming current levels of (for example) socioeconomic development, population and government capacity. Open circles, individual estimates; filled circles, mean across experts.

There is robust evidence and high agreement that socioeconomic development is the single most important covariate in cross-sectional studies and time series.⁵⁶⁴ Fearon and Laitin (2003) find a strong negative cross-country correlation between income per head and conflict propensity. Between 1945 and 2003, the chance of having a civil conflict for countries at different levels of income was as follows:

- 10th percentile in income per head (\$580): 72% chance
- A country at the median (\$2,671): 44% chance.
- A country at the 90th percentile (\$8,650): 2.7% chance.⁵⁶⁵

⁵⁶⁴ Mach et al., SI p2.

⁵⁶⁵ James D. Fearon and David D. Laitin, 'Ethnicity, Insurgency, and Civil War', *The American Political Science Review* 97, no. 1 (2003): 16.

In a future of low socioeconomic development and poor governance in low income countries, conflict is likely to be high anyway, regardless of climate change.

Future conflict predictions

There are interesting forecasting studies projecting levels of conflict forward based on different socioeconomic trends.

Hegre et al (2016) 'Forecasting Civil Conflict along the Shared Socioeconomic Pathways', *Environmental Research Letters*

Hegre et al - scholars at PRIO - note

"The conflict research community has identified a handful of robust country-level correlates of civil conflict, the three most powerful of which are a history of prior conflicts, a large population, and a low level of socioeconomic development"⁵⁶⁶

They quantify the size of these effects by analysing the correlates of major civil conflict (>1000 battle deaths).⁵⁶⁷

Population

The estimates for log population indicate that the odds of both minor and major conflict increase by 39% when population is increased by a factor of $e \approx 2.7$ – more populous countries have more frequent conflicts, but considerably less conflict per capita than smaller ones.

Income

Among countries currently at peace, increasing GDP per capita from e.g. USD 1,000 to USD 2,700 decreases the odds of major conflict by 20%.

Education

Controlling for the effect of GDP per capita, increasing YMHEP education by 0.1 (e.g., changing the proportion of the male population between 20 and 24 years from 30 to 40%) further reduces the odds of conflict by 20%.

Hegre et al (2016) models civil war risk (>1000 battle deaths) on the SSPs using their predictors of income per head, population size and history of conflict. Note this is excluding the effect of climate change - this is because (contra Burke et al who used a different methodology) they couldn't identify an effect on conflict in their statistical analysis.⁵⁶⁸ Burke et al would say that Hegre et al (2016) do not account for omitted variable bias.

⁵⁶⁶ Håvard Hegre et al., 'Forecasting Civil Conflict along the Shared Socioeconomic Pathways', *Environmental Research Letters* 11, no. 5 (April 2016): 054002, <https://doi.org/10.1088/1748-9326/11/5/054002>.

⁵⁶⁷ Hegre et al., SI.

⁵⁶⁸ "In a final test, we investigated a possible separate effect of temperature anomalies on conflict risk. This test revealed a weak and insignificant effect in the historical sample, and accounting for temperature anomalies did not improve the predictive performance of the model (section S5 in supplementary information)" Hegre et al., 'Forecasting Civil Conflict along the Shared Socioeconomic Pathways'.

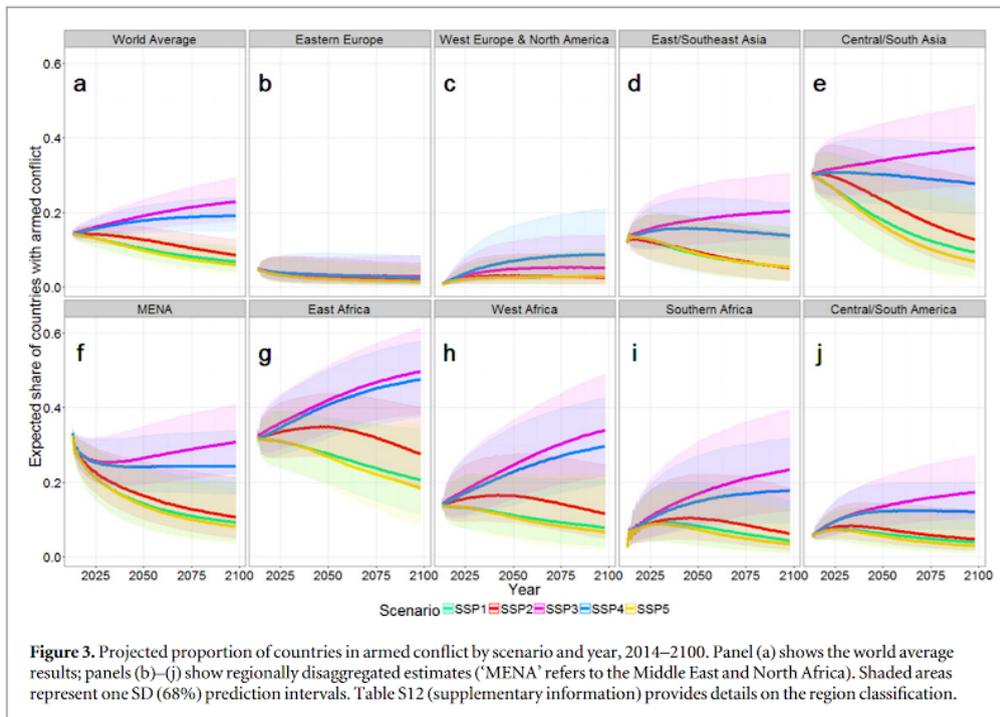


Figure 3. Projected proportion of countries in armed conflict by scenario and year, 2014–2100. Panel (a) shows the world average results; panels (b)–(j) show regionally disaggregated estimates ('MENA' refers to the Middle East and North Africa). Shaded areas represent one SD (68%) prediction intervals. Table S12 (supplementary information) provides details on the region classification.

As this shows, on all the SSPs except SSP3 and SSP4, conflict declines overall. Hegre et al (2013) produces similar findings of a decline in conflict, assuming that socioeconomic development follows past trends.⁵⁶⁹ In all regions, conflict declines the most on SSP5, the high growth future. Recall that SSP3 and SSP4 assume that there is no improvement in education in most/all countries. But they all also assume some convergence growth, which is at odds with the historical trend.

Low socioeconomic development and increased population in SSP3 and SSP4 leads to increased civil conflict in Africa by 200% by 2100. Note however that increasing population reduces conflict per capita.

Independently of climate change, if socioeconomic improvement is low, conflict will probably increase.

Witmer et al (2017) 'Subnational violent conflict forecasts for sub-Saharan Africa, 2015–65, using climate-sensitive models', *Journal of Peace Research*

Witmer et al (2017) projects future episodes of subnational political violence, which includes

“riots, protests, violence against civilians, and battles between rebel and government factions”⁵⁷⁰

⁵⁶⁹ Håvard Hegre et al., 'Predicting Armed Conflict, 2010–20501', *International Studies Quarterly* 57, no. 2 (1 June 2013): 250–70, <https://doi.org/10.1111/isqu.12007>.

⁵⁷⁰ Frank DW Witmer et al., 'Subnational Violent Conflict Forecasts for Sub-Saharan Africa, 2015–65, Using Climate-Sensitive Models', *Journal of Peace Research* 54, no. 2 (1 March 2017): 175–92, <https://doi.org/10.1177/0022343316682064>.

To generate forecasts of future violence, they compile historical data on the factors known to be the most important influences on violent conflict. In their approach, they control for key social and political variables, an approach the Berkeley economists think is mistaken.⁵⁷¹ The drivers of violence they consider are political rights, infant mortality, temperature and population growth.

Their results are as follows:

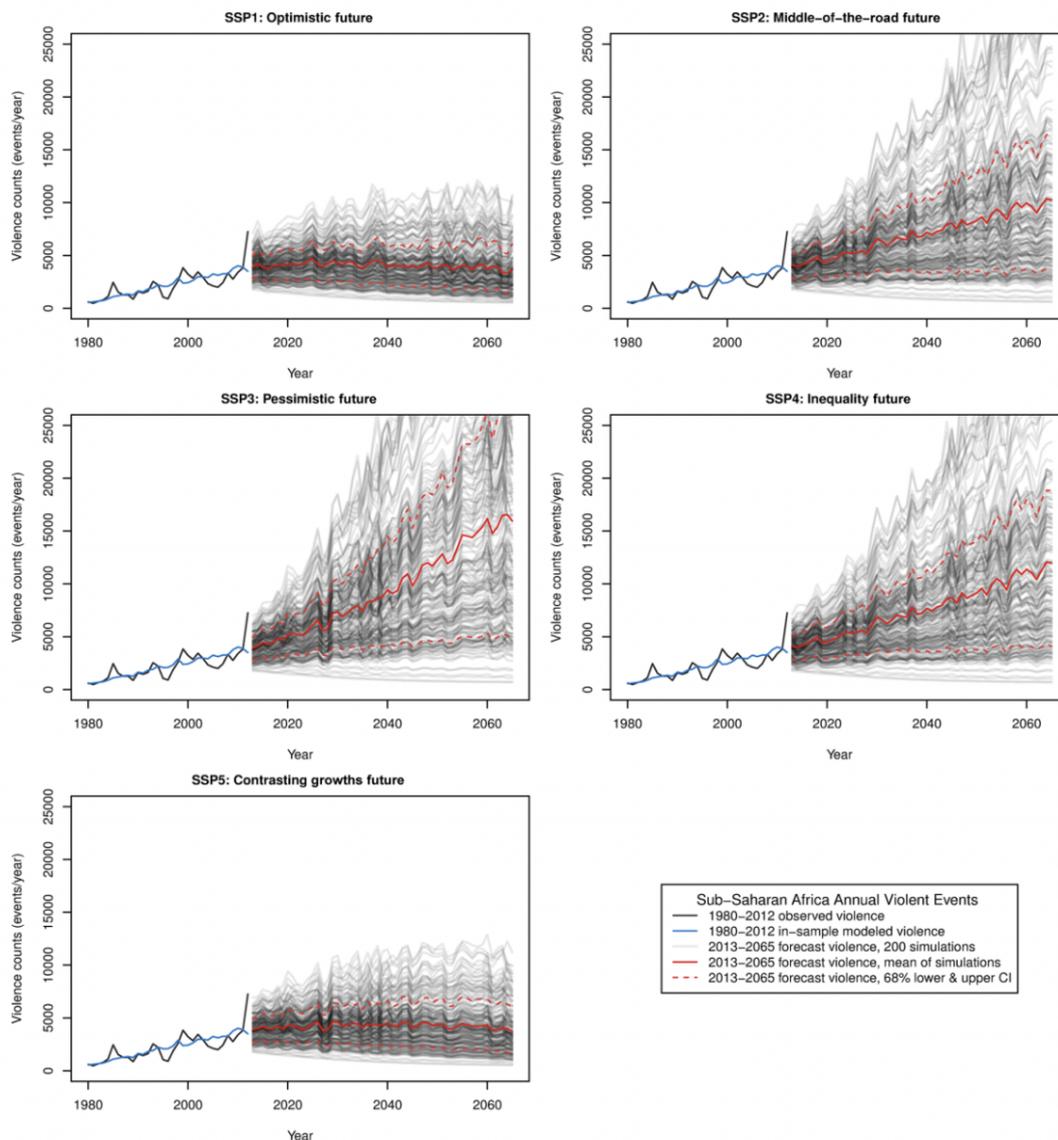


Figure 4. Forecast annual violent event counts for all of sub-Saharan Africa through 2065 for each of the defined shared socio-economic pathways (SSPs)

They find that political rights and population growth have a much greater influence on future conflict than climate change to 2065, which implies warming of around 3°C. For example, the effect size of poor governance is 20x larger than higher temperature.

⁵⁷¹ “In modeling and forecasting violence across all of sub-Saharan Africa at a 1 (degree) grid resolution, we explicitly consider the effects of temperature variability, while controlling for temporal reporting bias in the coding of our violent event data. Additionally, we account for a number of key social and political variables that have known associations with violence” Witmer et al.

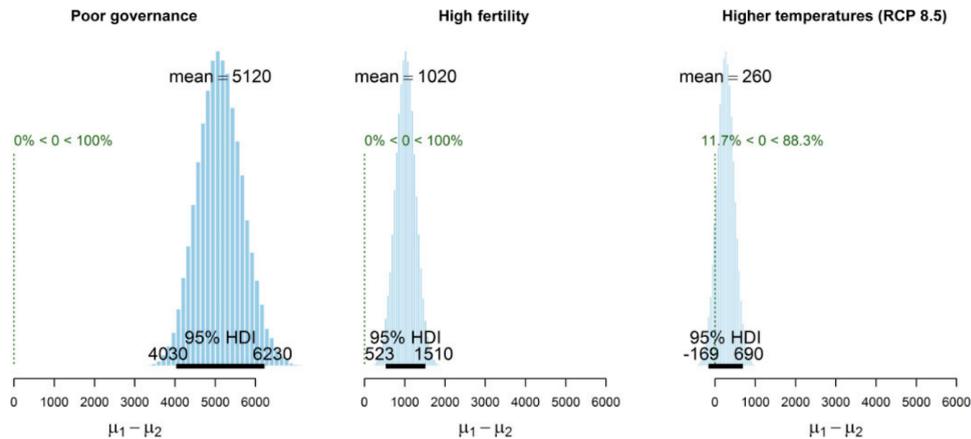


Figure 9. Posterior distributions of the difference in mean forecast violence for the period 2056–65 using SSP1 as the baseline forecast and comparing against poor governance, high fertility, and higher temperatures. The black bar at the bottom represents the highest density interval (HDI) and shows where 95% of the credible values fall

Each of these projection studies are dependent on choices about data, some of which the Berkeley economists would disagree with. Nonetheless, they do illustrate that other factors are considerably more important for civil conflict risk.

12.4. Concluding thoughts on climate change and civil conflict

I have argued that some prominent studies on the connection between climate change and civil conflict are flawed. This does not necessarily mean that there is no connection. The question of how climate change affects conflict is inherently difficult to study, though some insights can be drawn from the literature:

1. **Importance of agriculture.** The most plausible mechanism through which climate change could affect conflict is by harming agriculture. This will primarily be a problem for countries that remain poor and agrarian in the future. As I discussed in Chapter 5, given historical growth trends, there is a good chance that many countries in Sub-Saharan Africa and some in South Asia will remain poor and agrarian for much of the 21st Century. Climate change would be an additional stressor of conflict in those countries, but it seems unlikely to play a major role in countries with more advanced agricultural techniques and where food spending takes up a smaller share of average income.
2. **Socioeconomic development is key.** The first point implies that socioeconomic development will be the most important driver of trends in conflict in the 21st Century. If there is limited socioeconomic development, then civil conflict will remain a problem in Africa. If there is strong socioeconomic development, the prevalence of civil conflict seems likely to fall, as will the potential role of climate change as a stressor of conflict.

12.5. Climate, conflict and Great Power War

Unfortunately, the existing literature on climate change and conflict does not shed much light on the risk of future interstate conflict. Because interstate warfare is now so rare, the

literature has focused on intrastate civil conflicts. Moreover, among interstate conflicts, conflicts between the major powers pose by far the largest expected risk to humanity. This is because the major powers have far more destructive weaponry and have the capacity to alter the trajectory of humanity in other ways.

I have argued so far that the direct risks of climate change do not come close to causing permanent stagnation, civilisational collapse or human extinction. Economic models which add up the direct effects of climate change across different sectors find that the monetised costs of warming of 4°C are equivalent to a 5-10% of GDP in 2100 relative to a counterfactual without climate change (i.e. not relative to today). There is much less literature on the impacts of >5°C, so any judgments on the effects of that level of warming are much more speculative. Still, I find it difficult to see how even warming of 8°C could directly cause civilisational collapse or human extinction.

Consequently, if climate change could cause civilisational collapse or extinction, it would have to do so via its indirect effects, and the main candidate indirect effect that is large enough is a Great Power War. It is therefore crucial to assess the extent to which climate change is a stressor of Great Power War.

For people aiming to prioritise resources in order to have the greatest possible positive impact on the world, we first and foremost need to understand the size of the indirect risks of climate change *relative to* the total risk posed by other problem areas, such as nuclear security, AI safety and biorisk.

12.5.1. How would Great Power War affect the long-term future?

There is no universally accepted list of Great Powers, but it seems clear that given some combination of likely economic and military power, and the willingness to deploy it, over the course of the 21st Century, the list should include the US and other NATO members, China, Russia, and India. Here, I will also discuss Pakistan because it is a nuclear power which stands to be hit especially hard by climate change.

Great Power War would not itself cause a global catastrophe, but it is a risk factor for global catastrophe. It increases the risk of catastrophe from other causes, which includes nuclear war, engineered pathogens, and advanced AI. Toby Ord roughly estimates that if there were no Great Power War this century, existential risk would fall by around one tenth.⁵⁷² Since he estimates that existential risk this century is around 17%, Great Power War accounts for upwards of one percentage point of the existential risk this century.

In his [analysis](#) for the Forethought Foundation, Stephen Clare estimates that there is around a 1% to 2% chance of an existential disaster due to a Great Power War this century. This is almost entirely driven by the war and the build-up to the war causing states to develop and deploy dangerous weapons of mass destruction, including engineered bioweapons, misaligned AI systems, and currently unknown technologies. On Clare's model, a small fraction of the risk is due to nuclear war because it seems very unlikely that nuclear war could lead to human extinction.

⁵⁷² Toby Ord, *The Precipice: Existential Risk and the Future of Humanity* (Bloomsbury Publishing, 2020), 177.

Conflict dynamics in World War II led major powers to build atomic weapons. In the Cold War, countries hugely expanded their nuclear arsenals and built hydrogen bombs, and the Soviet Union built a huge bioweapons program. If tensions between the major powers increase again, major powers could once again race to develop even more destructive weapons.

Another important factor to consider is that Great Power War could be a key determinant of which country becomes the global hegemon this century, which could be the main determinant of which values get locked-in for the future. This could be comparably important to the direct extinction risks from war.

12.5.2. Theories of interstate war

Before we examine the causes of war, it is worth pausing to consider different theories of International Relations, as these have different implications for which causal drivers of war are most important.

The three main theories of the causes of interstate war are realism, liberalism and constructivism. All of these theories assume that states are unitary actors aiming to realise a grand strategy. An alternative view is public choice theory, which holds that states are not unitary actors with identifiable goals and strategies.

Realism

The realist family of theories hold that key actors in world politics are sovereign states that act to rationally advance their own security and power relative to other states in an anarchic international system. Anarchy, in this sense, refers to the absence of an international government authority to regulate disputes and enforce agreements between states and other actors. Despite their broad agreements, there are many different forms of realism, which sometimes disagree about the true causes of war and about the sign of the effect of different policies.⁵⁷³

Rational states?

The combination of international anarchy and states acting in their own interests might suggest that realists will be pessimistic about the prospects of global peace, and some scholars indeed contend that this is an implication of realism.⁵⁷⁴ However, in his classic article 'Rationalist explanations for war', James Fearon (1995) notes that the central puzzle

⁵⁷³ Jack S. Levy and William R. Thompson, *Causes of War*, 1st edition (Chichester, West Sussex, U.K.; Malden, MA: Wiley-Blackwell, 2010), Ch 2.

⁵⁷⁴ "intentions of other states, has enormous consequences. It induces insecurity and a continuous competition for power, which makes the international system inherently conflictual. Given omnipresent threats, political leaders tend to focus on short-term security needs and adopt worst-case thinking. They often utilize coercive threats to advance their interests, influence the adversary, and maintain their reputations. Anarchy does not automatically lead to war, but it creates a permissive environment for war by creating a system of insecurity, conflicts of interest, and international rivalries. Realists tend to have a pessimistic worldview, and they tend to be skeptical of grand schemes for creating and maintaining a peaceful international order." Levy and Thompson, *Causes of War*, p. 29.

of international relations is that wars are costly but nevertheless occur.⁵⁷⁵ It should usually be cheaper to negotiate a settlement than to go to war. So, if states are rational unitary actors, as posited by realism, war should be rare. Fearon outlines three potential causes of war, assuming that states are rational

1. Rational miscalculation due to private information about relative capabilities or resolve and incentives to misrepresent such information in order to deter attack.
2. Rationally led states may be unable to arrange a settlement that both would prefer to war due to commitment problems, situations in which mutually preferable bargains are unattainable because one or more states would have an incentive to renege on the terms. This is especially acute
3. There may be issue indivisibilities which do not admit compromise.

Fearon is more sceptical about the last possibility as a true rationalist explanation for war because rational states should be able to compensate one another for other issues that seem indivisible.

An alternative and reasonable alternative to the rationalist view, but one not compatible with realism as traditionally stated, is that sometimes states are not rational. This might be because leaders or voters are simply irrational about what would advance the national interest. Common sense and history suggest that this happens regularly.

Another possibility is the *unchecked interests* of leaders.⁵⁷⁶ The costs of war are the main incentive for peace, but when the people who decide on war aren't accountable to the others in their group, they can ignore some of the costs and agony of fighting. These leaders will take their group to war too frequently. Sometimes they expect to gain personally from conflict, and so they're enticed to start fights. Blattman comments that "unchecked rulers like these are one of the greatest drivers of conflict in history." This account has overlap with public choice theory, which I discuss below.

Liberalism

Like realists, liberal theories of international relations assume that states seek to maximise utility. But liberal theories are more hopeful about the scope for cooperation between rational states when mutual gains are possible, and about the role of international institutions in preventing conflict.⁵⁷⁷

Liberals make several substantive claims about the causes of war. First, like realists, they accept that international anarchy is a problem. However, they believe the problem can be reduced by the construction of international institutions like the EU and the UN.

Secondly, contra realists, they believe that not all states are equally likely to wage war. Democracies are less likely to wage war than autocracies, so the spread of democracy should reduce war.

⁵⁷⁵ James D. Fearon, 'Rationalist Explanations for War', *International Organization* 49, no. 03 (1995): 379–414.

⁵⁷⁶ Christopher Blattman, *Why We Fight: The Roots of War and the Paths to Peace* (New York: Viking, 2022), Ch. 2.

⁵⁷⁷ Levy and Thompson, *Causes of War*, p. 69ff; Cashman, *What Causes War?*, p. 170ff.

Third, economic interdependence through international trade can help to reduce the risk of war by increasing the opportunity cost of war and creating a constituency of special interests who stand to lose out from war and so would lobby against it.

Constructivism

Constructivists emphasise the role that identities and norms play in determining a state's actions.⁵⁷⁸ For example, a state's identity—perhaps as democratic, or developing, or hegemonic—will influence which norms it sees itself as following. Constructivist scholars emphasise that their approach allows them to consider a broader range of variables that bear on state decision-making.

From a constructivist point of view, states are not rational maximisers of the interests of their citizens. Norms and values determine the kinds of actions that states will pursue. For instance, human rights norms constrain the actions of states not because of considerations of power capabilities (and of consequent sanctions), but because human rights in part constitute the identity of states, especially states that are democratic.⁵⁷⁹

The structure of the international system - of states pursuing their material interests in a state of international anarchy - is a much weaker determinant of state behaviour than realists suggest. Constructivism mainly aims to explain war by explaining changes in individual beliefs, which are determined by shared norms and values. For instance, one explanation of why interstate war has declined is that there is now a widespread norm that offensive war is wrong.⁵⁸⁰

Public choice theory

According to public choice theory, contra realism, liberalism and (arguably) constructivism, states are not unitary rational actors that try to optimise a particular goal. Rather, state action is driven by people or factions who have a concentrated interest in the outcome, especially special interests in the defence establishment and weapons industries, and politicians trying to get re-elected. Decisions are made mostly on the basis of short-term political considerations, with a playing field tilted by concentrated interests.⁵⁸¹ This means that actions a state takes will often have limited consistency with any overall unified strategy, and will often foreseeably not help to enhance their power relative to other states, or to advance their citizens' interests.

From a public choice perspective, the most important drivers of conflict are politicians seeking to get re-elected or otherwise hold on to power, people in government defence departments, people in defence think tanks who are often government funded and regularly rotate between governments and think tanks, and finally weapons industries. Because the defence establishment wants to expand its budget, and military contractors want to make

⁵⁷⁸ Cashman, *What Causes War?*, p. 461ff.

⁵⁷⁹ Cashman, *What Causes War?*, p. 463.

⁵⁸⁰ Cashman, *What Causes War?*, p. 488-489.

⁵⁸¹ Richard Hanania, *Public Choice Theory and the Illusion of Grand Strategy: How Generals, Weapons Manufacturers, and Foreign Governments Shape American Foreign Policy*, 1st edition (Routledge, 2021), p. 4.

money selling weapons, each group has an incentive to inflate external threats and to argue against diplomatic solutions which might reduce defence spending.⁵⁸² Voters only weakly constrain the preferences of these concentrated interests because voters are uninformed about foreign policy and lack the incentives to learn more, meaning that they can be easily led by groups in the foreign policy establishment who are perceived to have expertise.⁵⁸³

On this view, military conflict will happen much more often than it would if states rationally pursued their self-interest; conflict is usually not justified by a rational analysis of its costs and benefits to the state or its citizens.

It should be stressed that the public choice theory is different to the view that states wage war in order to gain access to natural resources to benefit capitalists. Noam Chomsky is one prominent proponent of this analysis of the goals of US foreign policy: he argues that the US often goes to war to gain oil from other countries.⁵⁸⁴ But this theory assumes that the US is a unitary actor with a coherent strategy, which the public choice point of view denies.

12.5.3. How might climate change affect the risk of Great Power War?

I can think of four main mechanisms by which climate change might affect the risk of interstate war and, indirectly, war between great powers:

1. **Conflict over water resources:** Climate change changes the availability of water resources across the world, leading to conflict between countries over shared water resources.
2. **Economic costs:** Climate change imposes large economic costs, especially on low- and middle-income countries, which causes internal upheaval.
3. **Civil conflict:** Climate change increases the risk of civil conflict, which in turn increases the risk of interstate war.
4. **Mass migration:** Due to the economic and agricultural disruption in low- and middle-income countries, and to increased storms and flooding, there is a large increase in migration and displacement within and between countries, which increases political instability, increasing the risk of conflict.

All of these problems would be most severe in low and middle-income countries that are reliant on rainfed agriculture, and countries susceptible to riparian and coastal flooding.

Conflict over water resources

The possibility of 'water wars' sometimes comes up in discussions of how climate change might affect the risk of conflict. There currently are around 310 international river basins that are shared by 150 countries. They cover 47% of the world's land surface and are home to 52% of the world's population.⁵⁸⁵ Climate change will affect the supply of these resources,

⁵⁸² Hanania, *Public Choice Theory and the Illusion of Grand Strategy*, 53-54.

⁵⁸³ Hanania, *Public Choice Theory and the Illusion of Grand Strategy*, 45-46.

⁵⁸⁴ Chomsky, Noam, and David Barsamian, 2010, *Imperial Ambitions: Conversations on the Post-9/11 World*, Metropolitan Books .

⁵⁸⁵ Thomas Bernauer and Tobias Böhmelt, 'International Conflict and Cooperation over Freshwater Resources', *Nature Sustainability* 3, no. 5 (May 2020): 350-56, <https://doi.org/10.1038/s41893-020-0479-8>.

which could, it is argued, lead to increased interstate conflict. However, the evidence suggests that water resources are a weak driver of conflict.

Actual militarised conflict over water is extremely rare, and the overwhelming response to shared water resources has been cooperation. Wolf et al (2003) note that

“The only recorded incident of an outright war over water was 4500 years ago between two Mesopotamian city-states, Lagash and Umma, in the region we now call southern Iraq. Conversely, between the years 805 and 1984, countries signed more than 3600 water-related treaties, many showing great creativity in dealing with this critical resource”⁵⁸⁶

Wolf et al (2003) categorise all instances of water conflict and cooperation between 1948 and 2000 using the Basins at Risk (BAR) scale. The Basins at RISK scale ranges from +7 for unification into one nation, to -7 for formal war.

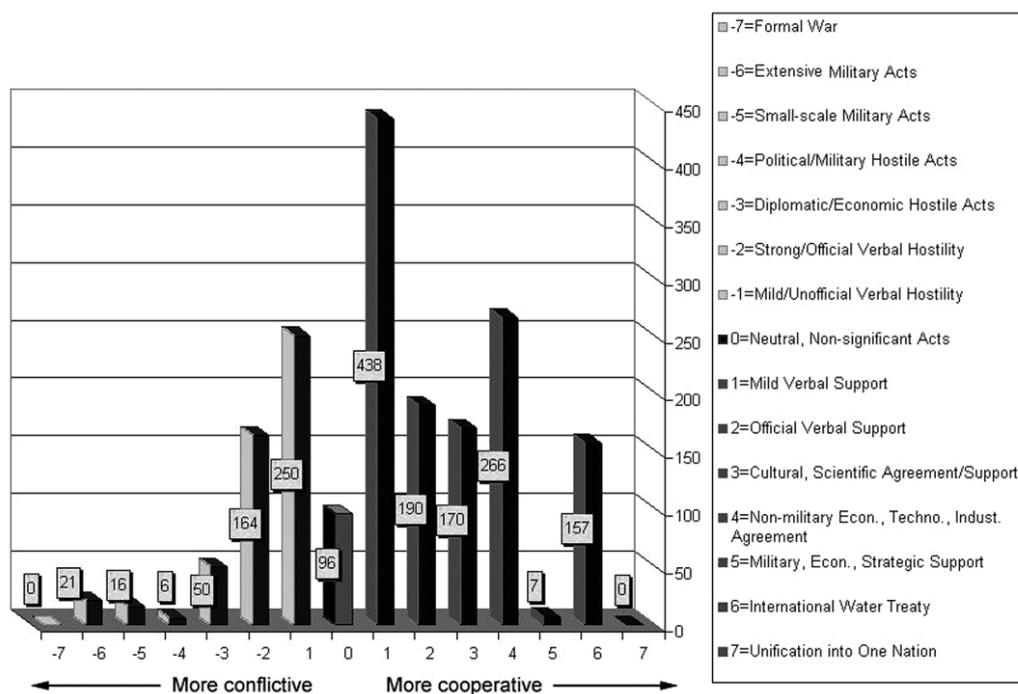


Fig. 3. Number of events by BAR Conflict/Cooperation Scale.

Source: Wolf, Aaron T., Shira B. Yoffe, and Mark Giordano. “International waters: identifying basins at risk.” *Water Policy* 5, no. 1 (2003): 29-60.

As this shows:

- There were no instances of formal war over water.
- Most interactions were mild, ranging between -2 and +4.
- The overwhelming pattern is one of cooperation, with international water treaties strongly outweighing conflict.

⁵⁸⁶ Wolf, Aaron T., Shira B. Yoffe, and Mark Giordano. “International waters: identifying basins at risk.” *Water Policy* 5, no. 1 (2003): 29-60.

The instances of conflict are not recent and largely occurred in the Middle East:

“Of the 37 cases of acute conflict, 30 are between Israel and one or other of its neighbors, violence which ended in 1970. Non-Middle East cases account for only five acute events.”⁵⁸⁷

The pattern of cooperation has been confirmed by more recent data.(see also de stefano) Bernauer and Böhmelt (2020) find that since 1948, states’ interaction over freshwater catchments is, on global average of all catchments and countries in a given year, characterised by more cooperation than conflict.

⁵⁸⁷ Wolf et al, “International waters: identifying basins at risk.”, 39.

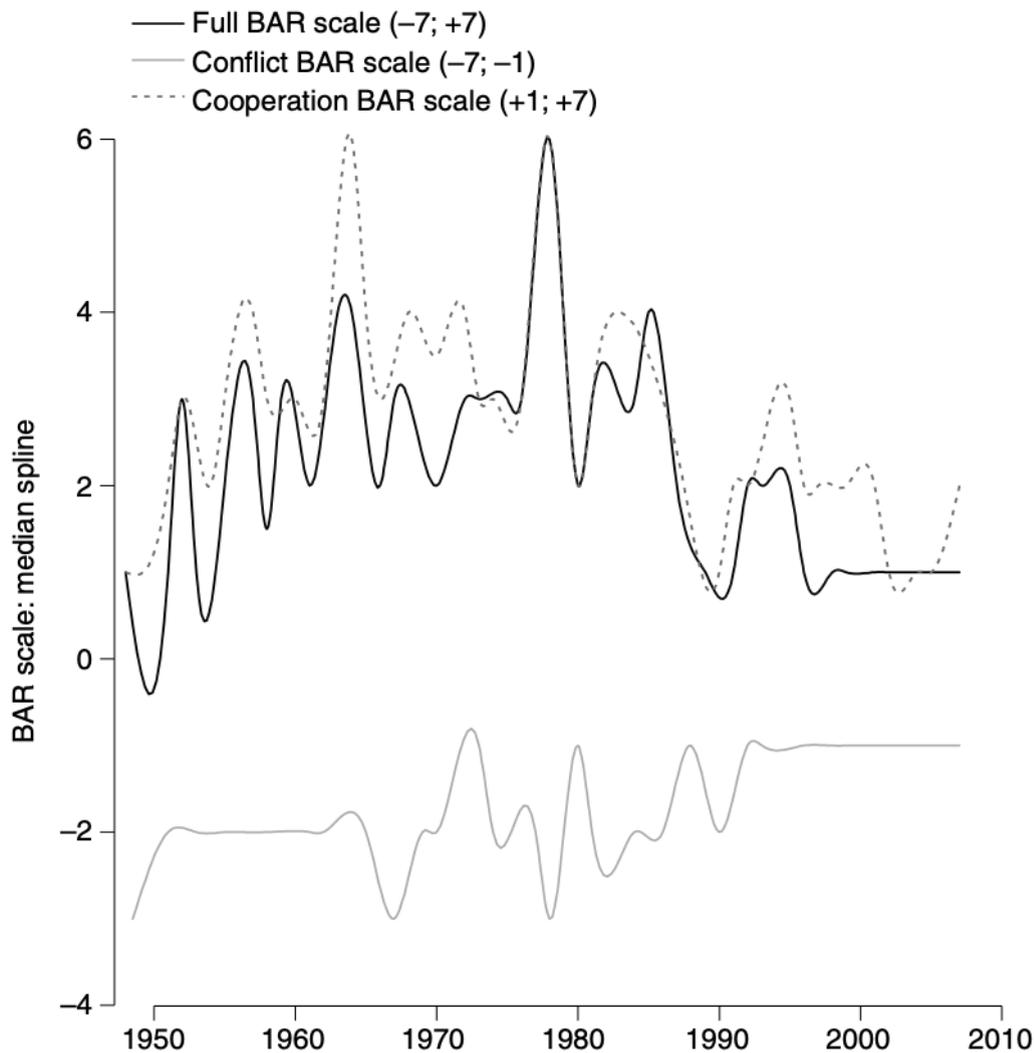


Fig. 1 | BAR scale, 1948–2008. The graph depicts three median splines across the BAR scale’s observation period for all possible BAR values, only cooperative (positive) ones, and only conflictive (negative) ones. The graph is based on data from www.transboundarywaters.science.oregonstate.edu.

Source: Thomas Bernauer and Tobias Böhmelt, ‘International Conflict and Cooperation over Freshwater Resources’, *Nature Sustainability* 3, no. 5 (May 2020): 350–56, <https://doi.org/10.1038/s41893-020-0479-8>.

Since 1950, the net global average Basins at Risk scale has been above zero, implying more cooperation than conflict.

There is evidence suggesting that cooperation over shared water resources increases the chance of broader cooperation between states. Ide et al (2018) found that “For the period 1956–2006, we find that a higher number of positive, water-related interactions in the

previous ten years makes a shift toward more peaceful interstate relations more likely. This is particularly the case for state pairs that are not in acute conflict with each other".⁵⁸⁸

Wolf et al (1998) notes that water acts as both an irritant and a unifier.

"The historical record shows that water disputes do get resolved, even among bitter enemies, and even as conflicts rage over other issues. Some of the most vociferous enemies around the world have negotiated water agreements or are in the process of doing so. The Mekong Committee has functioned since 1957, exchanging data throughout the Vietnam War. Secret "picnic table" talks have been held between Israel and Jordan since the unsuccessful Johnston negotiations of 1953–55, even as these riparians, until only recently, were still in a legal state of war. The Indus River Commission survived through two wars between India and Pakistan. And all ten Nile riparians are currently involved in negotiations over cooperative development of the basin."⁵⁸⁹

Overall, it is not clear what the sign of the effect of shared water resources is. If states do tend to cooperate over water, and that in turn leads to broader cooperation, shared water resources seem on balance to be a driver of cooperation rather than conflict.

From a theoretical point of view, it is also difficult to see why countries would go to war over water, rather than negotiate a settlement. As Wolf (1998) says "War over water seems neither strategically rational, hydrographically effective, nor economically viable."

"Water is neither a particularly costly commodity nor, given the financial resources to treat, store and deliver it, is it particularly scarce. Full-scale warfare, on the other hand, is tremendously expensive. A "water war" simply would not cost out.

This point was probably best made by the Israeli Defense Forces analyst responsible for longterm planning during the 1982 invasion of Lebanon. When asked whether water was a factor in decision-making, he noted, "Why go to war over water? For the price of one week's fighting, you could build five desalination plants. No loss of life, no international pressure, and a reliable supply you don't have to defend in hostile territory".⁵⁹⁰

This is one reason that the prospect of conflict over water is less plausible than conflict over more expensive commodities like oil and gas.

Economic costs

A commonly held belief is that states are more likely to go to war during times of economic distress. This might be because.⁵⁹¹

⁵⁸⁸ Tobias Ide and Adrien Detges, 'International Water Cooperation and Environmental Peacemaking', *Global Environmental Politics* 18, no. 4 (2018): 63–84.

⁵⁸⁹ Wolf, A.T. (1998). 'Conflict and cooperation along international waterways', *Water Policy*, 1(2), 251-265

⁵⁹⁰ Wolf (1998). 'Conflict and cooperation along international waterways'.

⁵⁹¹ Cashman, *What Causes War?*, 190ff.

- War may enable states greater access to markets and to resources.
- Leaders think war will stimulate the economy.
- Leaders may seek war to divert attention away from hard times.
- Leaders may be more willing to take risks during times of economic distress.

The evidence for this is mixed, and there is some evidence that the risk of interstate war is greatest during economic upswings. One reason for this might be that wars are more financially feasible in economic upswings.

A study of global economic cycles (called “long cycles”) and war from 1495 to 1975 by Joshua Goldstein finds a strong and consistent correlation between the *severity* of war and economic upswings. Although wars have occurred in roughly equal numbers throughout history in the upswing and downswing phases, the most severe wars have taken place in upswing phases. From 1495 until 1918 each peak in war severity occurred near the end of an upswing phase.⁵⁹²

Thompson (1982) used data on British, American, French and German business cycle phases and several types of war initiations from 1792 to 1973. He found that there was no association between war outbreak and the business.⁵⁹³

Cashman concludes that:

“Probably the most we can say about the business cycle is that it may play a role in the development of war, but its effects are less than clear. Some wars have broken out in hard times, others have occurred in good times. Neither economic weakness nor prosperity seem to prevent war.”⁵⁹⁴

Civil conflict and upheaval

A variety of studies have shown that there is a connection between internal upheaval and the risk of interstate war.⁵⁹⁵ Cashman (2009) summarises some of the quantitative evidence

“Later studies using more sophisticated methodology and larger databases have shown much more support for the linkage between internal and external conflict. For instance, Geller’s analysis of a variety of internal and external conflict variables for thirty-six states for the years 1959–68 shows that states with high levels of internal instability are more likely to be involved in external conflict than states with lower levels of instability. 5 Mansfield and Snyder find a very strong and statistically significant relationship between the presence of civil wars and interstate wars (but not world wars and not extra-systemic or colonial wars) for the period between 1816 and 1992.6 As part of their research they perform an analysis of the probability that a given dyad (a pair of countries) will be involved in an interstate war with each other. Their results indicate that if one country in the pair is undergoing civil war, this strongly increases the probability of war. 7 Another analysis by Houweling and

⁵⁹² Cashman, *What Causes War?*, 191.

⁵⁹³ William R. Thompson, ‘Phases of the Business Cycle and the Outbreak of War’, *International Studies Quarterly* 26, no. 2 (1 June 1982): 301–11, <https://doi.org/10.2307/2600653>.

⁵⁹⁴ Cashman, *What Causes War?*, 191.

⁵⁹⁵ Cashman, *What Causes War?*, Ch. 6.

Siccama demonstrated that civil wars and interstate wars cluster in both time and space, implying that the presence of one increases the probability of the other. 8 Cashman and Robinson's case studies of seven wars in the last one hundred years (though based on an unscientific sample) found that in five of the six wars between equals, domestic political variables helped to explain the initiation and escalation of conflict. In these cases, one of the participant states had significant domestic political instability and/or ruling elites perceived themselves to be politically vulnerable. (The cases are Iran prior to the Iran-Iraq War, Pakistan prior to the 1971 war with India, Japan in the 1930s and 1940s, Ethiopia in 1998, and Germany and Austria in 1914.)⁵⁹⁶

The theoretical explanations for this are as follows:⁵⁹⁷

- **Diversionsary:** leaders could use war to divert public attention away from internal turmoil. This is especially risky for enduring rivals.
 - Cashman (2009) comments that “Ultimately, it is not entirely clear whether the theory of diversionsary wars—applicable to a wide variety of countries and representing a general pattern—has much basis in reality, and if so what kinds of countries are most likely to use diversionsary force and under what conditions. The support for the theory is certainly not very robust. Diversionsary war appears to be a distinct path to war, but one that is very little traveled”⁵⁹⁸
- The “kick ‘em while they’re down” theory: predatory states could exploit states who are weakened by internal upheaval.
 - Cashman (2009) notes that “The kick-’em-while-they’re-down theory certainly is not an explanation that can be widely applied; only a small percentage of all wars fall into this category.”⁵⁹⁹
- **Revolutions** alter the balance of threats and power between states and revolutionary states often face a hostile political environment.
 - There is good evidence that violent revolution increases the risk of war. Of nine violent revolutions between 1789 and 1979 highlighted by Stephen Walt,, six led to wars within five years; the other three were near-misses that resulted in militarized disputes short of full-scale war.⁶⁰⁰
- **The internationalisation of civil conflicts** in part due to external intervention.
 - Gleditsch, Salehyan and Schultz (2008) found that states experiencing civil wars are substantially more likely to engage in military disputes than others.⁶⁰¹ They find that the presence of a civil war increased the probability of a militarised interstate dispute (which falls short of war) in a dyad by between 50 percent and 80 percent over the baseline predicted probability depending on the civil war data set used. They argue that “intervention, externalization, and unintended spillover effects are important sources of international

⁵⁹⁶ Cashman, *What Causes War?*, 200.

⁵⁹⁷ Cashman, *What Causes War?*, Ch. 6.

⁵⁹⁸ Cashman, *What Causes War?*, 210.

⁵⁹⁹ Cashman, *What Causes War?*, 214.

⁶⁰⁰ Cashman, *What Causes War?*, 215-218.

⁶⁰¹ Kristian Skrede Gleditsch, Idean Salehyan, and Kenneth Schultz, ‘Fighting at Home, Fighting Abroad: How Civil Wars Lead to International Disputes’, *Journal of Conflict Resolution* 52, no. 4 (2008): 479–506.

friction". The increased probability of a militarised interstate dispute due to the existence of a civil war was roughly equal to the decreased probability of a dispute due to joint democracy.

Mass migration, displacement and conflict

I discussed the potential effects of climate change on migration and displacement in the last Chapter. To recap:

1. The character of future environment-related population movement
 - a. Most environment-related population movement is short-term, internal and temporary
 - b. Most environment-related population movement is likely to occur in low- and middle-income countries.
2. Displacement
 - a. Weather-related displacement is likely to predominantly occur in Asia, mainly due to storms and flooding.
 - b. Weather-related coastal displacement for 4°C is likely to be strongly outweighed by migration to coastal regions in Asia.
 - c. Evidence on displacement from river flooding is weaker, but my rough estimates suggest that warming of 4°C would increase overall displacement from conflict and environmental disasters by around 10% relative to today.
3. Migration
 - a. It is much harder to quantify the overall effect of climate change on migration
 - b. Studies have found that environmental change can both increase and decrease migration.
 - c. The overall evidence suggests that climate change will probably increase migration, though the size of the effect is very unclear.
 - d. The estimates we do have suggest that climate change would increase internal migration by around 10%, though this estimate may be substantially biased in either direction.

In short, the overall connection between climate change and population movement is unclear and disputed. The connection between population movement and violent conflict is also disputed.

The mechanism

There are two main mechanisms proposed in the literature by which migration or displacement increases the risk of civil (as opposed to) conflict. These are:

1. The Malthusian idea that economic migrants increase conflict over resources.⁶⁰²

⁶⁰² "The influx of large numbers of "environmental migrants" is likely to burden economic and resource bases in the receiving areas, thus promoting contests over scarce resources." Vally Koubi, 'Climate Change and Conflict', *Annual Review of Political Science* 22, no. 1 (11 May 2019): 343–60, <https://doi.org/10.1146/annurev-polisci-050317-070830>.

2. Population movement could inflame ethnic or cultural tensions with the host population.⁶⁰³

The plausibility of the mechanisms depends in part on the nature of the migration or displacement. Migration due to a short-term climatic event is less likely to cause civil conflict than migration due to a long-term climatic event such as a drought. Brzosky and Fröhlich summarise the risks associated with different types of environmental migration:

Table 2. Risk factors for violent conflict of different types of climate-related migration.

Type	Conflicts over interests	Identity conflicts	Conflict processes	Likelihood of violent conflict
Eco-economic migrants	Conflict over employment opportunities	Movements into areas with strong and adverse identities	‘Otherness’ discourses in receiving regions	Generally low except in regions with strong identities and ‘otherness’ discourses
Climate disaster refugees	Conflict over resources where they are constrained and humanitarian assistance is scarce or not forthcoming	Large-scale movement into areas occupied by hostile identity group	Climate disasters affecting groups involved in violent conflict	Low, because of widespread availability of humanitarian assistance and low capabilities to organize violence
Permanent climate refugees	Conflict over resources, employment opportunities	Movements into areas with strong and adverse identities	Discourses over economic competition and ‘otherness’ in receiving regions	Depending on intensity of conflicts over interests and identities
Climate-affected migrants	Conflict over resources	Conflict over identities	Dependent on absence, presence of institutions for conflict management	Depending on conflict intensity and absence, presence of institutions for conflict management

Source: Michael Brzoska and Christiane Fröhlich, ‘Climate Change, Migration and Violent Conflict: Vulnerabilities, Pathways and Adaptation Strategies’, *Migration and Development* 5, no. 2 (3 May 2016): 190–210, <https://doi.org/10.1080/21632324.2015.1022973>.

Salehyan (2008) discusses two mechanisms by which international refugees could increase the risk of interstate conflict:⁶⁰⁴

⁶⁰³ “In addition, environmental migration could lead to conflict by stirring ethnic tensions that arise when migrants and residents belong to different ethnocultural groups and the arrival of newcomers upsets an unstable ethnopolitical balance” Vally Koubi, ‘Climate Change and Conflict’.

⁶⁰⁴ Idean Salehyan, ‘The Externalities of Civil Strife: Refugees as a Source of International Conflict’, *American Journal of Political Science* 52, no. 4 (2008): 787–801, <https://doi.org/10.1111/j.1540-5907.2008.00343.x>.

1. Refugees constitute a negative externality borne by receiving states, and these states may launch military actions to seal their borders, threaten sending regimes with violence, and even invade the sending state to prevent further flows.
2. Refugee-sending states may violate the sovereign territory of their neighbours in order to attack political and/or ethnic rivals that have fled across the border as well as punish the states that harbour them.

The size of the effect on conflict

Unsurprisingly, there is disagreement about the size and sign of the effect of migration and displacement on the risk of both civil and interstate war. The IPCC's Sixth Assessment Report has the following findings on the connection between climate-related population movement and conflict:⁶⁰⁵

1. There is some evidence of an association between climate-related rural-urban migration and civil unrest such as riots and protests in urban areas.
2. There is evidence that climate-related internal migration has been associated with the prolongation of conflict in migrant-receiving areas.
3. There has been no association established between international migration and conflict.

The evidence we have suggests that displaced people are less likely to cause civil conflict than climate-affected migrants because displaced people are likely to be the victims of storms and floods, which mainly cause short-term disruption.

However, the findings of the literature on the connection between migration and conflict are mixed.⁶⁰⁶

Although not mentioned by the Sixth Assessment Report, Salehyan (2008) explores the connection between cross-border flows of refugees and 'militarised interstate disputes'.⁶⁰⁷ Militarised interstate disputes are events that involve at least two states in which there was a

⁶⁰⁵ Source: IPCC, *Impacts*, Sixth Assessment Report, Ch 7, sec. 7.2.7.

⁶⁰⁶ "While these and other models are plausible, they need to be substantiated through empirical evidence. Empirical support so far is scant. The available empirical evidence, while generally supporting the importance of mediating factors identified in the theoretical literature on the causal pathway from climate change to migration to violent conflict, crucially points to the under-complexity of existing models. As mentioned above, effects of climate change on local livelihoods differ in several dimensions. Furthermore, only in extreme cases are decisions to migrate shaped by livelihood conditions in regions affected by climate change, and even then, factors unrelated to climate change may be important, such as the availability of external humanitarian assistance or existing social networks. In most cases, the decision to migrate will be shaped by many factors beyond changes in resource availability in a location, including conditions in potential receiving regions but also group, family and gender relations. Similarly, case studies point to the complexity of the relationship between population movements and violent conflict. Even where violent conflict has often been seen as most likely to erupt, such as in the Sahel zone, where farmers and herders compete for scarce resources, the evidence is mixed, and it remains difficult to explain reality with linear models linking migration to violent conflict (Benjaminsen, Alinon, Buhaug, & Buseeth, 2012; Raleigh et al., 2010). Moreover, case studies seemingly contradict each other, making generalized statements difficult." Michael Brzoska and Christiane Fröhlich, 'Climate Change, Migration and Violent Conflict: Vulnerabilities, Pathways and Adaptation Strategies', *Migration and Development* 5, no. 2 (3 May 2016): 190–210, <https://doi.org/10.1080/21632324.2015.1022973>.

⁶⁰⁷ Salehyan, 'The Externalities of Civil Strife'.

threat, display, or use of military force. Thus a military interstate dispute need not actually involve war.

Before presenting his quantitative analysis, Salehyan (2008) outlines two illustrative case studies: the Indian invasion of East Pakistan (now Bangladesh) in 1971, and the Rwandan invasion of Zaire in 1996. In both cases, cross-border refugees seemed a plausible cause of or contributing factor to interstate conflict. Salehyan (2008) constructs a regression model which attempts to isolate the effect of refugee inflows on the probability of conflict between pairs of countries that either share a boundary or one is a major power.⁶⁰⁸

The model produces the following results

TABLE 2 Predicted Probabilities

	Prediction	Std. Err.	% Change
Baseline*	0.012	0.002	
100k refugees in initiator	0.023	0.006	96.55
100k refugees from target	0.023	0.004	90.34
Civil war in initiator	0.021	0.003	77.46
Civil war in target	0.017	0.002	40.46
Joint democracy	0.006	0.001	-53.16

*Baseline: no refugees, no civil wars, neither democratic, neither transitional, contiguous dyad, all other variables at mean.

100,000 refugees is roughly the mean amount of refugees in the dataset for all nonzero observations. So, the baseline risk for a pair of countries which share a border, are not democracies and have no civil wars is around 1.2%. An increase in refugees by 100,000 increases the risk of interstate dispute by around one percentage point. This is larger than the effect of both countries becoming democracies, which reduces the risk of dispute by 0.6 percentage points.

Quantifying the overall effect of migration and displacement on interstate war is very difficult given how rare interstate war now is.

12.5.4. The main drivers of interstate war

All of the theories outlined above have different views about the most important drivers of the risk of interstate war. Wars are usually due to a complex mixture of factors that interact with

⁶⁰⁸ "Because many pairs of states are unlikely to be involved in military conflicts with one another (e.g., Peru and Tanzania), following convention, I restrict the analysis to politically relevant dyads where the states are either contiguous or involve at least one major power" Michael Brzoska and Christiane Fröhlich, 'Climate Change, Migration and Violent Conflict: Vulnerabilities, Pathways and Adaptation Strategies', *Migration and Development* 5, no. 2 (3 May 2016): 190–210, <https://doi.org/10.1080/21632324.2015.1022973>.

one another, and the factors will vary from case to case.⁶⁰⁹ Research has however identified certain drivers and triggers that increase the risk of war. Below, I highlight those most relevant to the risk of Great Power War.⁶¹⁰ All of these factors are correlated and endogenous. Each of the theories outlined above will have different views on how important these factors are, or indeed whether they matter at all.

War usually involves contiguous neighbours

For obvious reasons, neighbouring countries are more likely to fight one another. If war does not involve contiguous neighbours, then it is likely to involve a major power.

A large proportion of wars involve enduring or strategic rivals

Rivals are pairs of states who see each other as threatening.⁶¹¹ Relations between these states tend to be fraught with hostility and they tend to have had repeated instances of diplomatic and military conflict over time. Obvious examples include India and Pakistan, and Israel and Egypt.

Security dilemma and conflict spiral

A conflict spiral occurs when states respond to hostility with increased hostility. As the level of hostility escalates, conflict spirals may develop that eventually lead to war. This is closely related to the security dilemma, which is, as Jervis puts it, that most means of self-protection simultaneously menace others.⁶¹² In a situation of international anarchy, where there is no central authority to restrain states from violence, each state is compelled to provide for its own security. However, the security policies of states are interdependent; greater security for one state may mean relatively less security for others.

According to Levy and Thompson, “a large array of scientific studies provide evidence to support a conflict spiral theory of international conflict”.⁶¹³ Cashman claims that “most wars are preceded by militarized disputes or crises that involve escalatory behavior preceding the outbreak of war that looks like a conflict spiral—though the temporal length of the spiral may vary considerably”.⁶¹⁴

Unresolved territorial disputes

The most likely cause of a conflict spiral is an unresolved territorial dispute. This is especially true for neighbouring countries.

Misperceptions about relative power

Misperceptions about the opponent’s actions, intentions and capabilities and hence about the degree of threat to one’s own security may create the conditions necessary for the onset of a crisis.⁶¹⁵ In fact, misperceptions by national leaders have frequently been cited as the

⁶⁰⁹ Cashman, *What Causes War?*, 478.

⁶¹⁰ These are factors identified in Cashman, *What Causes War?*, Ch. 13; Levy and Thompson, *Causes of War*.

⁶¹¹ Cashman, *What Causes War?*, 252.

⁶¹² Levy and Thompson, *Causes of War*, 290.

⁶¹³ Levy and Thompson, *Causes of War*, 289

⁶¹⁴ Cashman, *What Causes War?*, 479.

⁶¹⁵ Cashman, *What Causes War?*, 492.

immediate cause of war. World War I, Vietnam, Korea, the Middle East Wars of 1967 and 1973, the first Gulf War, and the Iraq War of 2003 are typically seen as examples. Cashman and Robinson look at seven cases of interstate war in the last 100 years and find that “in almost no case did the leaders in the initiating country, operating under what we have called the fog of prewar, accurately perceive the situation in which they found themselves.”⁶¹⁶

According to Cashman:

“Once the crisis begins, these misperceptions may accelerate and exacerbate the level of tension. Particularly important are the combination of an over-perception of the rival’s hostility, treachery, and threat, coupled with an under-perception of both the capabilities of the rival and the amount of risk involved. Unwarranted confidence in the ability to compel one’s opponents to back down short of war, or of one’s ability to defeat the adversary with little cost if war does begin, seem to be an important part of the picture.”⁶¹⁷

Misperceptions are not necessarily irrational. They can arise because each side has private information about relative capabilities, and has incentives to misrepresent such information.

Power transitions

One of the most important and oft-discussed potential drivers of Great Power War is power transition.⁶¹⁸ According to power transition theory, the most dangerous and war-prone situation is one in which a state that is rising and dissatisfied with the status quo begins to approach the strength of the leading state in the system and threatens to surpass it in power. Conflict is more likely to occur when the competing powers are roughly equal in power, so there is a window in the transition process when conflict is especially likely.

Power transitions could be driven by differential rates of economic growth or military accumulation, or new strategic alliances by the rising power.

There are several reasons to think that conflict is more likely during power transitions. Firstly, there is a greater risk of the security dilemma and conflict spirals. As the rising power gains more economic and military power, they may prompt the hegemon to further increase their own military power. This could in turn lead to escalating conflict

Secondly, at times of great change in power, there will also be great uncertainty about the relative power of each actor. Misperceptions about relative power are a key driver of war.

Thirdly, commitment problems are severe during power transitions. The rising power could promise the current hegemon that they will not take advantage of their newfound status once they are on top, but there is no way to enforce such a promise. This creates incentives for the current hegemon to pre-emptively strike before it loses hegemonic status.

⁶¹⁶ Cashman, *What Causes War?*, 99.

⁶¹⁷ Cashman, *What Causes War?*, 492.

⁶¹⁸ “Serious great-power crises have in the past been most likely to occur during periods of transition in the international system (or in regional subsystems) where there are significant shifts in the balance of capabilities, especially between the dominant power in the system and its major rival(s)—but also between any set of rivals.” Cashman, *What Causes War?*, 485.

12.5.4. The influence of climate change on the risk of Great Power War

In this section, I will outline the main drivers of potential conflicts between the major powers this century

The US and China

As of 5th August 2022, the community forecasting platform Metaculus puts the chance of a [war between the US and China](#) before 2035 at 16%.⁶¹⁹

By far the most discussed drivers of the risk of war between China and the US are: (1) the possibility of power transition - that China will surpass the US as the global hegemon in the next few decades,⁶²⁰ and (2) flashpoints about contested territories, including Taiwan and the Senkaku Islands. This assessment is widely shared in US foreign policy circles.⁶²¹

It is difficult to see how climate change could have much of an effect on these factors. On both the economic models that I think are most plausible and the pessimistic study by Burke et al (2015b), climate change of 4°C would have only small effects on the economic prospects of the US and China. There is little reason to think that climate change would have much effect on the absolute or relative economic performance of the US and China.

It is also clear that climate change has only negligible relevance to the most pressing territorial disputes relevant to US-China relations. China's views about Taiwan being part of One China, for instance, are at best tangentially related to climate change.

From a public choice point of view, much of the risk of war between China and the US is driven by special interests who have a vested interest in threat inflation, defence spending and bellicosity. For example, Richard Hanania, a proponent of the public choice theory, argues that defending Taiwan does not serve American interests. Nonetheless, concentrated interests have an incentive to inflate the importance of this conflict. The risk is, therefore, driven by the domestic political economy of the US and the fact that China has a strong ideological commitment to Taiwan being part of One China. Again, it is difficult to see how climate change is relevant to either of these factors.

The US and Russia

As of 5th August 2022, forecasters on Metaculus put the chance of a [war between the US and Russia](#) before 2050 at 25%.

⁶¹⁹ Note that for these forecasts, the resolution criteria may be different to what some people intuitively think of as war.

⁶²⁰ Graham T. Allison, *Destined for War: Can America and China Escape Thucydides's Trap?* (Houghton Mifflin Harcourt, 2018).

⁶²¹ "Destined for War, published in 2017 and named a notable book of the year by the New York Times, Financial Times, and Times of London, is the most influential work pushing the Thucydides Trap as a way to understand current international politics. The book has garnered praise from the likes of former CIA director David Petraeus, former Secretary of State Henry Kissinger, former Senator Sam Nunn, and former Secretaries of Defense Ash Carter and William Cohen. It also inspired the Harvard Thucydides's Trap Project, an ongoing effort to expand on and facilitate discussion about Allison's findings, created by the scholar himself. Cited nearly 800 times as of this writing, perhaps no international relations book of the last decade has had as much impact" Richard Hanania, 'Graham Allison and the Thucydides Trap Myth', *Strategic Studies Quarterly* 15, no. 4 (2021): 13–24.

Although there was no Great Power War after 1945, the US and the Soviet Union came close to nuclear war on several occasions. The drivers of these near-misses usually centred around (1) attempts to alter the balance of power, for example during the Cuban Missile Crisis, and (2) accidents with nuclear weapons, especially mistaken beliefs that the opposing side had launched a nuclear first strike.⁶²² The risk of such accidents was particularly acute during times of crisis.

After the end of the Cold War, tensions between the US and Russia declined, but they have increased again recently with Russia's invasion of Ukraine, which substantially increased the relative risk of a nuclear exchange between NATO powers and Russia.

There are two main causal explanations for Russia's invasion. One is that the war is a consequence of the US and its allies refusing to rule out Ukraine eventually joining NATO, which, from the Russian point of view, imposes intolerable security risks on Russia.⁶²³ Proponents of this theory have argued that NATO promised Russia not to expand eastwards after the fall of the Berlin Wall, but they broke that promise. So, Russia had incentives to strike before Ukraine had the chance to join NATO.

Another view is that the conflict is a product of Russian nationalistic aggression that would have occurred even if NATO membership was ruled out for Ukraine.⁶²⁴ The argument is that Putin does not see Ukraine as an independent sovereign state and so has violated the post-World War II norm against territorial conquest.

These two different theories will likely give different accounts of the future drivers of the risk of US-Russia war. On the former account, most of the risk stems from potential eastward expansion of NATO to include countries like Sweden and Finland. Finland shares a border with Russia and the two countries have a history of conflict.

The main alternative view is that future conflicts will be driven by increasing Russian nationalism and unprovoked military aggression against its neighbours, which Putin uses as a tool to cement his domestic dominance.

It is difficult to see how climate change is particularly relevant on either view of the future drivers of US-Russia conflict risk.

India and Pakistan

As of 5th August 2022, forecasters on Metaculus put the chance of a [violent conflict between India and Pakistan](#) killing at least 100 people by 2050 at 61%. There are no forecasts on the risk of all-out war between the two countries.

⁶²² Toby Ord, *The Precipice: Existential Risk and the Future of Humanity* (Bloomsbury Publishing, 2020), chap. 4.

⁶²³ John Mearsheimer, 'The Causes and Consequences of the Ukraine Crisis', *The National Interest*, June 2022, <https://nationalinterest.org/feature/causes-and-consequences-ukraine-crisis-203182>.

⁶²⁴ Tanisha Fazal, 'The Return of Conquest? Why the Future of Global Order Hinges on Ukraine', *Foreign Affairs*, June 2022, <https://www.foreignaffairs.com/articles/ukraine/2022-04-06/ukraine-russia-war-return-conquest>.

The rivalry contains many of the ingredients for conflict. First, the two countries have a history of violent conflict. They have engaged in major conflict three times over the past 60 years (1948, 1965, 1971), and a more limited war over Kargil in 1999.⁶²⁵ Tensions remain high today.

“In February 2019, an attack on a convoy of Indian paramilitary forces in Indian-controlled Kashmir killed at least forty soldiers. The attack, claimed by Pakistani militant group Jaish-e-Mohammad, was the deadliest attack in Kashmir in three decades. Two weeks later, India claimed to have conducted air strikes targeting a terrorist training camp inside Pakistani territory. Pakistan retaliated a day later with air strikes in Indian-administered Kashmir. The exchange escalated into an aerial engagement, during which Pakistan shot down two Indian military aircraft and captured an Indian pilot; the pilot was released two days later.”⁶²⁶

Second, the countries share a border. Thirdly, the countries have an unresolved territorial dispute over Kashmir, which was the cause of two of three major Indo-Pakistan wars in 1947 and 1965, as well as the smaller 1999 war.⁶²⁷ Fourthly, while both countries are currently democracies, Pakistan only became a democracy in 2003 and its status seems quite fragile. If either country becomes an autocracy, the risk of war would increase.

The character of the tensions between India and Pakistan changed dramatically once each country acquired nuclear weapons - India in 1974 and Pakistan in 1998. The stakes of conflict are now much higher, but conflict is also less likely.

It is much more plausible that climate change could play a role in future conflicts between India and Pakistan than for the other potential conflicts I have considered. Around [40% of the labour force](#) in each country is in agriculture, a highly climate-exposed sector. Both countries are at low latitudes and so will be hit especially hard by rising heat stress, especially for 4°C of warming. Thus, there is more scope for the climate-related drivers of war - economic disruption, migration and civil conflict - to play a role.

Water wars

I argued above that conflicts over water resources are a weak driver of interstate conflict. In the case of India and Pakistan, the Indus River Commission which manages water resources between the two countries, survived through two wars between India and Pakistan.

Economic performance

I argued above that the connection between economic disruption and war seems weak. But independent of that, different climate-economy models produce very different estimates of the costs of climate change for India and Pakistan. One approach is bottom-up: to add up the costs of climate change sector by sector using the impacts literature. As I discussed in Chapter 10, two such examples are Kompas et al (2018) and Takakura et al (2021). Alternatively, one can take a top-down approach. The most pessimistic study, Burke et al

⁶²⁵ Levy and Thompson, *Causes of War*, 4.

⁶²⁶ See [Global Conflict Tracker](#)

⁶²⁷ See [Global Conflict Tracker](#)

(2015a) calculates the effects of climate change by inferring from the past effects on economic output of interannual weather variation.

The tables below show the effects on GDP per capita, according to these different models, first for 2.5°C warming, which is most likely on current policy, and then for 4°C, which has a roughly 5% chance on current policy.⁶²⁸

Warming vs pre-industrial	2.5°C					
Damage	GDP/person in 2100 vs. no climate change			GDP/person in 2100 vs. today (SSP4)		
Study	Kompas	Takakura	Burke	Kompas	Takakura	Burke
India	-8%	-4%	-50%	+1595%	+1668%	+821%
Pakistan	-5%	-4%	-50%	+2275%	+2300%	+1150%

Warming vs pre-industrial	4°C					
Damage	GDP/person in 2100 vs. no climate change			GDP/person in 2100 vs. today (SSP4)		
Study	Kompas	Takakura	Burke	Kompas	Takakura	Burke
India	-15%	-8%	-80%	+1466%	+1595%	+268%
Pakistan	-9%	-8%	-80%	+2175%	+2200%	+400%

Source: Calculations are in [this google sheet](#)

As this shows, on the bottom-up models, the effects of climate change are fairly small relative to overall economic growth in the 21st Century. If these models are correct, climate change will be a relatively small driver of economic disruption in these countries. Alternatively, on the Burke et al (2015a) estimate, the effects of climate change on the future economic prospects of India and Pakistan are huge.

I explained in Chapter 10 why in my view, the bottom-up approach is more plausible, but probably understate the costs of climate change. I could imagine the damage estimates for 4°C being wrong by a factor of 2-4. However, the probability of 4°C now seems to be 5%, if not lower. Even if you think the costs of 4°C would be 40% of GDP, the expected economic costs would be around 2% of GDP, making climate change a small factor relative to the other determinants of war.

Moreover, there is limited evidence that poor economic performance is a major driver of the risk of interstate war.

⁶²⁸ Note that Takakura et al (2021) measure the monetised cost of non-market damages, so their damage estimate has a different meaning to the other two studies. Note also that Kompas et al and Takakura et al measure the effect on GDP not GDP per capita. However, in their models, climate change only affects the numerator of GDP per capita, so this does not affect the results.

Mass migration and civil conflict

The other two climate-related drivers I discussed in 12.5.3 were mass migration and civil conflict. These two factors are related because one way that migration might be thought to influence the risk of conflict is by influencing the risk of civil conflict.

India will be one of the countries worst-hit by displacement from coastal flooding, with on the order of tens of thousands of extra people displaced per year, with adaptation, for 1 metre of sea level rise relative to 30cm (see Chapter 7). The evidence suggests that coastal displacements will be outweighed by migration to coastal regions. However, due to progress in emissions, 75cm of sea level rise rather than 1 metre now seems like the most likely scenario. So, this is likely something of an overestimate of the increase in coastal displacements due to climate change.

Data on displacements from river flooding is worse, and it is hard to give an overall figure, but my own guess based on the limited data is that river flood displacements would be comparable to coastal displacements.

In 2020, in India 3.9 million people were displaced by weather-related and geophysical disasters, and in Pakistan the number was 829,000.⁶²⁹ Thus, the increased displacements from storms and flooding from climate change in these countries seem likely to account for a relatively small fraction of overall displacements.

It is much harder to predict how climate change will affect voluntary migration within India and Pakistan. Since economic factors are the main driver of migration, the size of the effect depends mostly on the size of the effect climate change will have on economic factors in India and Pakistan, which I discussed above. On the models that in my view are more plausible, climate change is a relatively small factor relative to the other determinants of economic growth in those countries, especially given recent progress in emissions.

The final part of the migration causal chain is 'migration => conflict'. As I discussed in Chapter 11 and earlier in this chapter, there is great uncertainty about the size of the effect of migration on conflict. The two mechanisms discussed in the literature are:

- (a) Cross-border refugees => interstate conflict
- (b) Migration => civil conflict => interstate conflict

On (a), there is also evidence that large flows of refugees across borders increase the risk of conflict between states and indeed that was a driver of one previous conflict between India and Pakistan in 1971, when refugees fled what was then East Pakistan (now Bangladesh) into India following conflict between the Pakistan army and Bengali separatists.⁶³⁰ However, most weather-related migration is typically short-term, internal and temporary. The vast majority of weather-related migrants in Pakistan or India are likely to stay in-country. Cross-border India-Pakistan migration is likely to be a small fraction of overall climate-related migration.

⁶²⁹ IDMC, *Global Report on Internal Displacement 2021*, p. 51.

⁶³⁰ Salehyan, 'The Externalities of Civil Strife', 792-793.

On (b), there is some evidence that mass migration increases civil unrest and protest in urban areas and prolongs ongoing civil conflicts. This might indirectly increase the risk of political upheaval in India and Pakistan and therefore the risk of interstate war.

Other climate impacts might also increase the risk of civil conflict, as discussed earlier in this chapter. If you take the Burke et al (2015b) estimate of the effect of climate change on conflict at face value, then 2°C would increase the risk of civil conflict by around 50%, while 4°C would increase it by around 100%. However, this estimate is mainly informed by data from Africa and I think it is a substantial overestimate, probably by a factor of two or more.

Below, I will construct a rough model that will bring all of these factors together to estimate the size of the effect of climate change on conflict.

India and China

As of 5th August 2022, forecasters on Metaculus put the probability of a [war between India and China](#) before 2035 at 20%.

The relationship has many of the ingredients necessary for conflict. Firstly, India is democratic and China is authoritarian. Secondly, the two countries have a history of conflict over disputed territory on their shared border. India and China went to war in 1962 over disputed border territories around Aksai Chin.⁶³¹ Indian forces were soundly defeated, 7,000 men having been killed or captured. The border dispute continues to this day.⁶³² Fighting in 2020 resulted in the deaths of both Indian and Chinese soldiers, and shots were fired in September 2020 for the first time in 45 years.⁶³³

India is also a member of the Quad group of countries (along with US, Japan and Australia) and it is unclear what role they will play in defending Taiwan if China ever does invade.⁶³⁴

Climate change seems like a weaker driver of this potential conflict than the India-Pakistan conflict. The main reason for this is that China is now relatively wealthy and so will be able to adapt to climate damages, and models suggest that it will not be hit especially hard by plausible levels of climate change. Moreover, climate change seems unlikely to affect the balance of power between the two countries: China's GDP per capita is already tenfold higher than India's, and China is experiencing much higher rates of economic growth. The only plausible mechanism of impact is if climate change causes civil unrest in India, which increases the risk of interstate war.

As discussed above, climate change might also drive migration and displacement in each country, but the vast majority of climate-related migrants are likely to remain in their home

⁶³¹ Britannica, [Sino-Indian War 1962](#).

⁶³² 'India-China Dispute: The Border Row Explained in 400 Words', *BBC News*, 25 January 2021, sec. Asia, <https://www.bbc.com/news/world-asia-53062484>.

⁶³³ 'First Time in 45 Years, Shots Fired along LAC as Troops Foil China's Bid to Take a Key Height', *The Indian Express*, 9 September 2020, 45, <https://indianexpress.com/article/india/first-time-in-45-years-shots-fired-along-lac-as-troops-foil-chinas-bid-to-take-a-key-height-6588658/>.

⁶³⁴ Huynh Tam Sang and Trang Huynh Le, 'How Far Could the Quad Support Taiwan?', *The Diplomat*, August 2022.

country. In the absence of political repression, it is difficult to see why there would be substantial climate-induced cross-border refugee movements between India and China.

12.5.5. Which conflicts would do the most damage?

Some of the conflicts outlined above have the capacity to be far more destructive than others, and have the potential to have a disproportionate impact on how the future will go. How potentially important these conflicts are depends in large part on the overall level of the risk posed by the different direct risks (AI, biorisk and nuclear war) of which Great Power War is a stressor.

In *The Precipice*, Toby Ord argues that AI and biorisk account for the majority of direct existential risk in the next 100 years.⁶³⁵

- AI accounts for 10 percentage points
- Biorisk accounts for 3pp
- Unforeseen anthropogenic risks account for 3pp
- Other anthropogenic risks account for 3pp
- Nuclear war accounts for 0.1pp.

For Ord, most of the existential risk we face this century will be driven by emerging technologies with destructive power, especially AI and biotechnology. Similarly, on Clare's model, the indirect risk of Great Power War is driven by the prospect that the competing powers would develop and deploy novel weapons of mass destruction.

AI

Insofar as Great Power War is a stressor of the risk from AI, this must almost entirely be driven by the risk of war between the US (and its allies) and China. The US and China account for around 60% of all AI publications and 90% of highly cited AI publications.⁶³⁶ The US, China and the EU account for 84% of private investment in AI.⁶³⁷ It seems highly plausible that transformative AI systems will be developed either by American companies or the Chinese government.

One concern is that increasing tension between the US and China could cause each of these countries to race to develop advanced AI systems that would enable them to gain strategic dominance over the other. This could lead to an existential catastrophe either due to accidents as each country skimps on safety, or deliberate misuse by the controlling government.

Since climate change is a weak lever on the risk of US-China conflict, it is also a weak lever on the risk of AI-induced existential catastrophe.

⁶³⁵ Toby Ord, *The Precipice: Existential Risk and the Future of Humanity* (Bloomsbury Publishing, 2020), 167.

⁶³⁶ Ashwin Acharya Brian Dunn, 'Comparing U.S. and Chinese Contributions to High-Impact AI Research', CSET, Jan 2022, Fig. 1, <https://cset.georgetown.edu/publication/comparing-u-s-and-chinese-contributions-to-high-impact-ai-research/>

⁶³⁷ This data is from [Statista](https://www.statista.com).

Biorisk

The capacity to invest in bioweapons is more evenly spread across the world. Despite having relatively low GDP per capita, the Soviet Union developed Biopreparat, the world's largest bioweapons program, which started in 1974 and at its height employed more than 9,000 scientists to weaponise various diseases, including smallpox and anthrax.⁶³⁸ It is, therefore, possible for relatively poorer countries to develop dangerous bioweapons programs if they are in a conflict situation and perceive there to be security benefits from doing so.

The development of such weapons seems strategically irrational, but such programs can nevertheless exist.

The extent to which Great Power War might be a stressor of biorisk depends on the nature of future biorisk. The existential risk from biotechnology might come from:

1. Accidental release from public health-focused research in university or government labs
2. Accidental or deliberate release from a state bioweapons program
3. Deliberate release by terrorists

If most of the risk will come from public health-focused research in university or government labs, then the US, China and the EU account for most of the risk as they have the greatest capacity to fund such research, and are its main funders today. The vast majority of Biosafety Level-4 labs, which handle the world's most dangerous pathogens, are in the EU, China and the US.⁶³⁹ This seems likely to remain true in the coming decades, given the economic power of these regions. It is not clear whether increased conflict would increase the risk of this kind of research. If it does, then climate change is a weak lever on this mechanism because it is a weak lever on the risk of US-China conflict.

It is difficult to know how state bioweapons programs might develop over the coming century, especially in the face of increasing tension between great powers. The *capacity* to build state bioweapons programs is likely to be concentrated in the US, the EU and China given their economic power and dominance of current biotechnology research.

However, it is unclear whether they will have the greatest *willingness* to develop bioweapons programs. Russia has shown the greatest historical willingness to develop a substantial bioweapons program, and the US believes that it still has one today.⁶⁴⁰ Insofar as Russia drives much of the future risk of accidental or deliberate release, the main conflict-related driver of this risk is the conflict between NATO powers and Russia, which is Russia's main security concern. As I argued above, climate change seems like a weak lever on this conflict.

⁶³⁸ Toby Ord, *The Precipice*, 202, 132.

⁶³⁹ Filippa Lentzos and Gregory Koblenz, 'Fifty-nine labs around world handle the deadliest pathogens – only a quarter score high on safety', 14 June 2021, King's College, <https://www.kcl.ac.uk/fifty-nine-labs-around-world-handle-the-deadliest-pathogens-only-a-quarter-score-high-on-safety>

⁶⁴⁰ "The United States assesses that the Russian Federation (Russia) maintains an offensive BW program and is in violation of its obligation under Articles I and II of the BWC. The issue of compliance by Russia with the BWC has been of concern for many years." [US state department](#).

As technology develops, bioweapons will become increasingly accessible to countries with more limited economic power, and some of these countries may be willing to develop bioweapons programs. For example, there is some evidence that China and Pakistan are collaborating on a program which will develop novel biological weapons.⁶⁴¹ So, it may well be that increased international tension could drive developing countries to develop dangerous bioweapons programs.

In the future, the capabilities of terrorists to develop bioweapons may also be spread across the world. Tension between the Great Powers is a weaker lever on this than on state bioweapons programs, but may nevertheless have some effect, as it might reduce the ability of states to cooperate to reduce bioterror.

Nuclear weapons

At present, the US and Russia hold the [vast majority](#) of the world's nuclear warheads. Assuming that this remains true in the future, the US and Russia account for the majority of the risk from nuclear war in the next 100 years. However, it is possible that if tensions rise, other states will also start to build up their nuclear arsenals and so the picture may change in the future. Indeed, China looks set to increase its arsenal in the future.⁶⁴²

Other technological risks

China, the EU and the US have the greatest capacity to invest not just in AI but other potentially destructive technologies. China, the US and the EU account for the vast majority of global spending on R&D.⁶⁴³ Thus, these regions plausibly account for most of the risk from novel weapons of mass destruction this century.

12.5.6. Quantifying the indirect risks of climate change

It is extremely difficult to provide reliable quantitative estimates of the risk of Great Power War caused by climate change. Nonetheless, I have built a model that attempts to put some numbers on the arguments I have discussed so far. I think this is valuable for several reasons. Firstly, it clarifies the cruxes of disagreements and allows focused discussion on those cruxes. Secondly, it allows us to prioritise different problems. If we do not quantify, we will still have judgments about how important different considerations are. Models make these considerations precise.

The downside of quantitative models is that they can cause false precision and anchor readers, even if the model is not good and has not been subject to scrutiny. Many of the considerations I have discussed are very difficult to quantify because there is essentially no literature on them.

⁶⁴¹ 'China Entered Covert Deal with Pakistan Military for Bio-Warfare Capabilities against India, Western Countries: Report', *The Times of India*, 24 July 2020, <https://timesofindia.indiatimes.com/india/china-entered-covert-deal-with-pakistan-military-for-bio-warfare-capabilities-against-india-western-countries-report/articleshow/77139556.cms>.

⁶⁴² Tong Zhao, 'What's Driving China's Nuclear Buildup?', Aug 2021, Carnegie Endowment for Peace.

⁶⁴³ This data is summarised by Wikipedia [here](#).

With those caveats in my mind, my model is [here](#). Given the discussion in this Chapter, I assume that all of the risk stems from the India v Pakistan and India v China conflicts, and in turn that most of the risk of existential catastrophe stems from AI, biorisk and currently unforeseen technological risks. My best guess estimate is that the risk of ‘climate change => great power war => existential catastrophe’ is on the order of 1 in 100,000. My pessimistic estimate puts the figure at 1 in 1,000. The pessimistic estimate makes the following assumptions, among others:

- Conflict between India and Pakistan accounts for:
 - 10% of the risk from ‘Great Power => AI’.
 - 30% of ‘Great Power War => biorisk’
 - 20% of ‘Great Power War => other anthropogenic risk’
- Climate change accounts for a third of civil conflict risk in India and Pakistan, and civil conflict accounts for 20% of the total risk of war between India and Pakistan this century.

12.5.7. Overall judgement on climate change and Great Power War

Overall, climate change seems likely to impose burdens on developing countries that will increase the risk of instability and civil conflict. This in turn is a plausible stressor or risk between major powers: worlds with greater instability and fewer economic opportunities broadly seem likely to be ones with more interstate war. Thus, I think the overall sign of the effect of climate change on interstate conflict is clear.

One highly controversial but commonly discussed example of this kind of effect is the Syrian Civil War. Some scholars argue that climate change-induced multi-year droughts played a major role in driving this conflict, though this is heavily disputed.⁶⁴⁴

During the conflict, the US and Russia provided military support to the opposing sides, and carried out air strikes against anti-Assad and pro-Assad forces, respectively. US commandos destroyed a force of Russian mercenaries with strong links to the Russian government,⁶⁴⁵ though the risk of a hot war between the US and Russia still seemed small.⁶⁴⁶

Even if climate change did not in fact play a major role in the Syrian Civil War, it is certainly plausible that it could contribute to increasing civil conflict in the future. Thus, this does illustrate one way that a domestic conflict could spill over to increase tensions between major powers.

⁶⁴⁴ “A small set of published studies has argued inconclusively over the role of drought in causing the Syrian civil war (Gleick, 34 2014);(Kelley et al., 2015);(Selby et al., 2017) [also 16.2.3.9].”

“Migration from drought-stricken areas to local urban centres has been used to suggest a climate trigger for the Syrian conflict (e.g.(Ash and Obradovich, 2020)). However, this link has 22 been strongly contested by research that contextualizes the drought in wider political economic approaches and existing migration patterns (De Châtel, 2014);(Fröhlich, 2016);(Selby, 2019) [16.2.3.9].” IPCC, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Sixth Assessment Report, 2022, Ch. 7, sec. 7.2.3

⁶⁴⁵ Thomas Gibbons-Neff, ‘How a 4-Hour Battle Between Russian Mercenaries and U.S. Commandos Unfolded in Syria’, *The New York Times*, 24 May 2018, sec. World, <https://www.nytimes.com/2018/05/24/world/middleeast/american-commandos-russian-mercenaries-syria.html>.

⁶⁴⁶ There is an overview of the Syrian Civil War in the Encyclopaedia Britannica [here](#).

However, I also think it is true that climate change is a relatively weak lever on future interstate conflict. It is difficult to come up with clear causal stories which explain why levels of climate change that now seem plausible would make states go to war. This is especially true of many of the most potentially destructive conflicts, between the US and Russia and the US and China. But it also seems true for conflicts between India and Pakistan, countries with nuclear weapons, a history of conflict, and which will be relatively badly affected by climate change. Although they will be badly affected, the levels of climate change that now seem likely have relatively little impact relative to the other determinants of conflict risk between those countries.

Moreover, working on climate change seems like a highly indirect way to work on reducing the risk of Great Power War. It seems to be more impactful to work directly on the most threatening conflicts and on the levers that are widely recognised to drive them, which include:

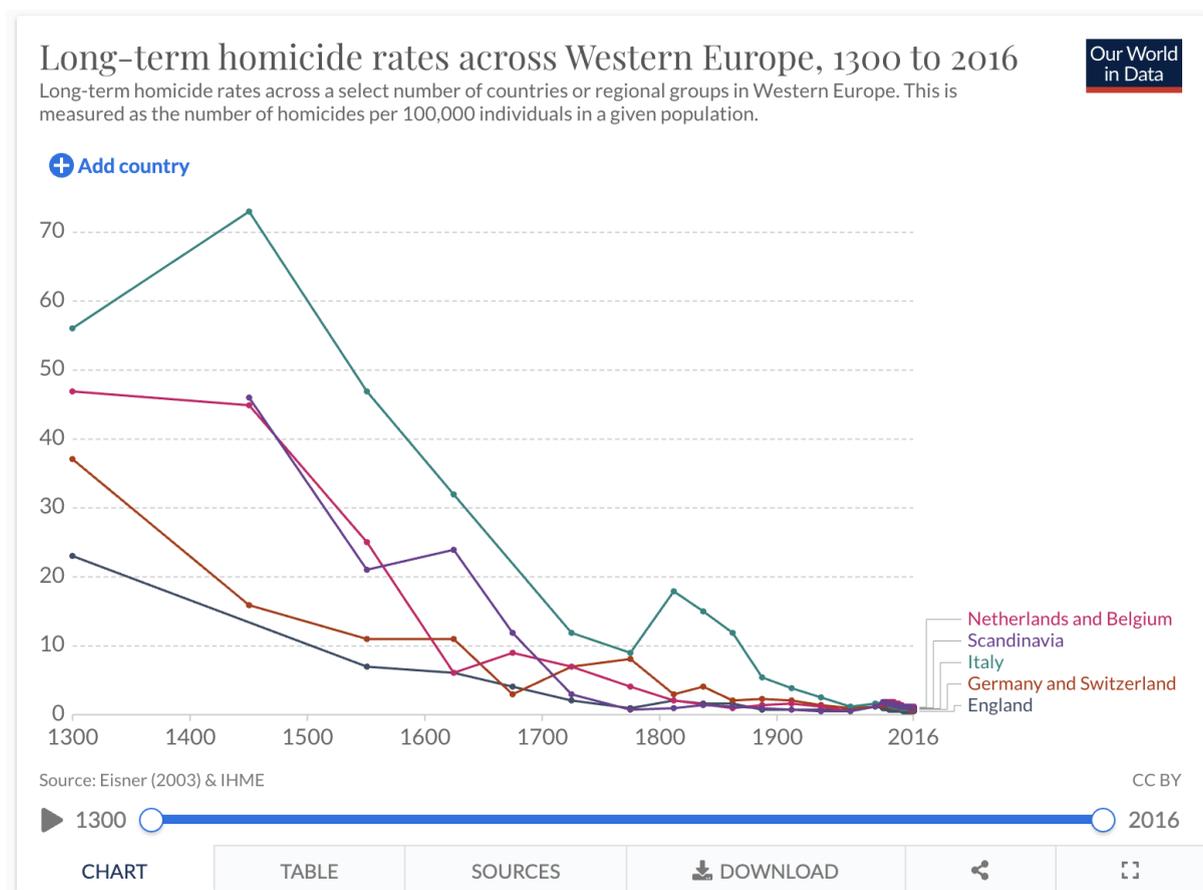
- Finding the right balance between deterrence and the conflict spiral, especially with respect to disputed territories
- Avoiding race dynamics for military technology and weapons of mass destruction
- Strengthening international management of emerging weapons of mass destruction
- Making a clear national strategy to manage power transitions
- Etc.

13. Interpersonal conflict, aggression and violence

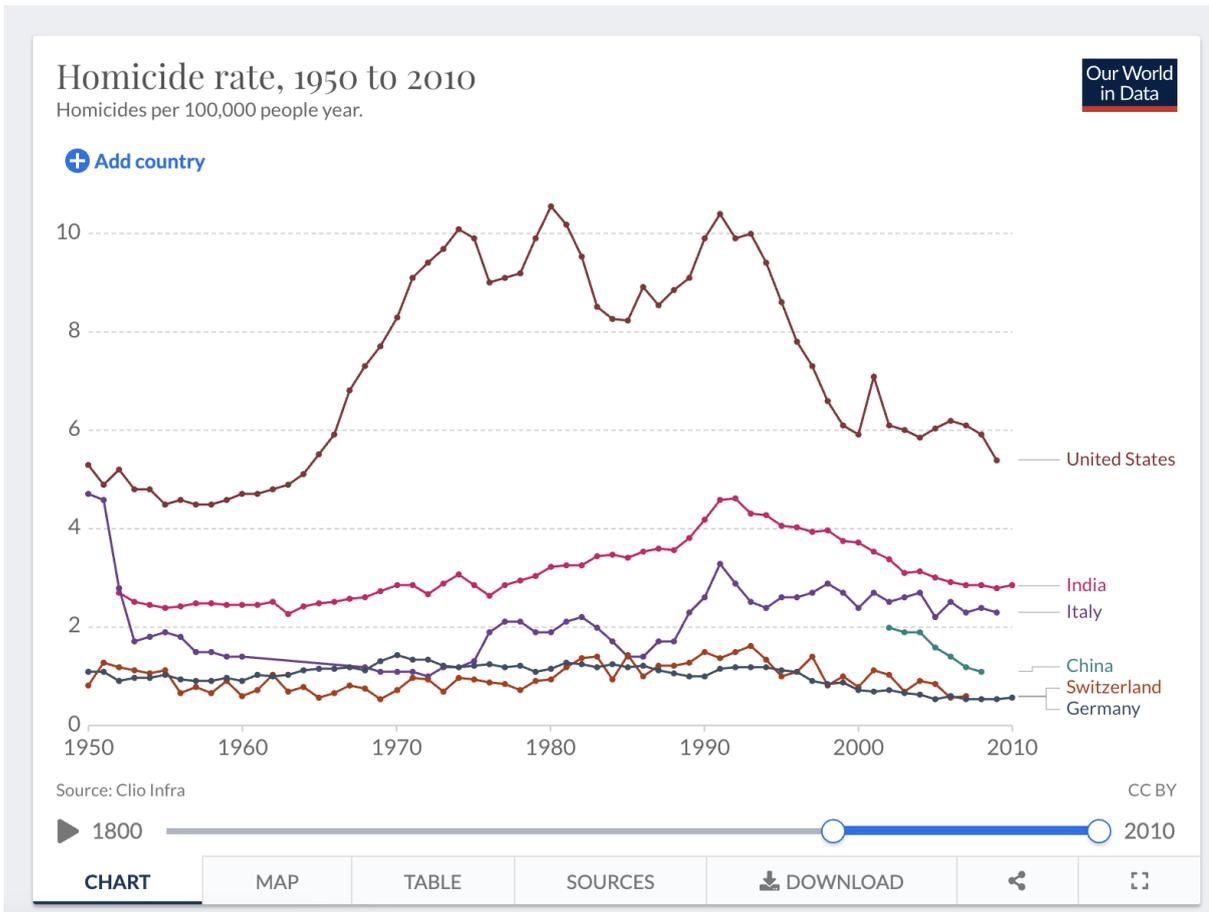
There is some evidence in the literature that climate change could affect levels of aggression and violence at the person level. One might worry about this because climate change has global effects and raising impulsivity, aggression and violent inclination across the whole global population seems robustly bad.

13.1. Context, base rates and trends

Homicide rates have been declining in Western Europe and the US for hundreds of years.

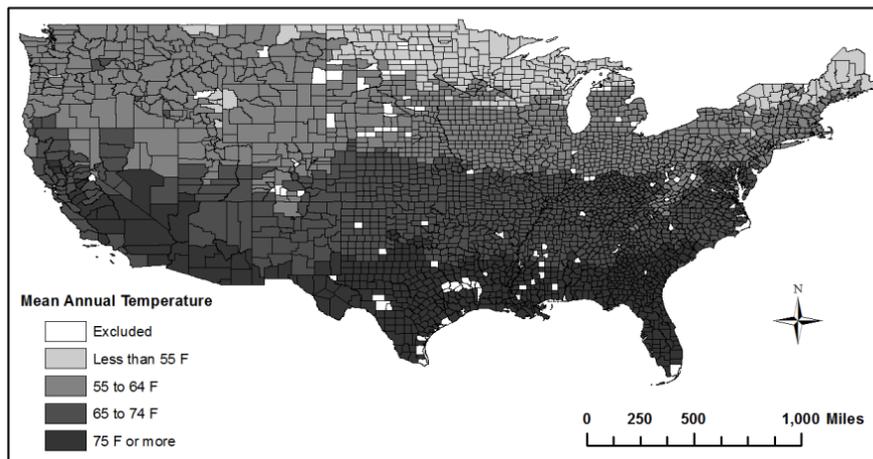


Rates of violent crimes vary substantially on short timescales. For example between 1990 and 2000, the homicide rate in the US dropped by 40%.

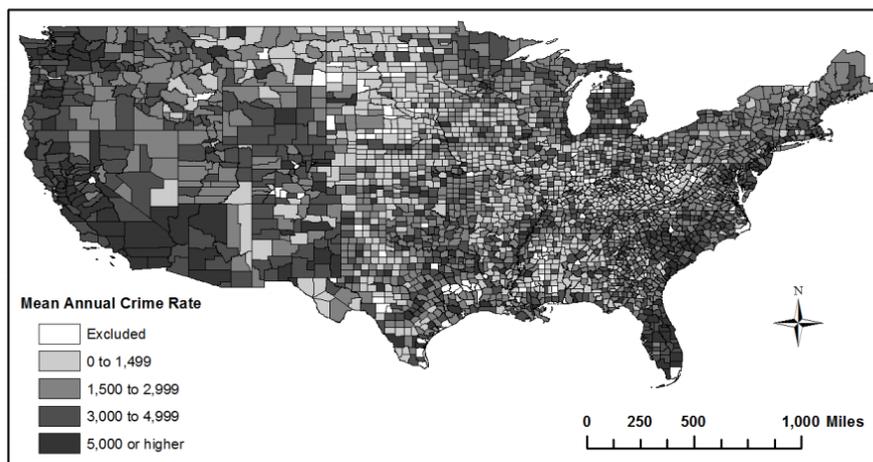


13.2. Cross-sectional relationship

The cross-sectional relationship between temperature and rates of violent crime within jurisdictions and between jurisdictions is at best not very strong. For example, within the US, this is the relationship between temperature and violent crime - just from eyeballing, the relationship seems quite weak:



(a) Mean Annual Maximum Daily Temperature (F)



(b) Annual Crime Rate per 100,000 Persons (All Crimes)

Figure 1: Map of the Study Region

Note: Both panels show maps of all in-sample counties in the United States. The top panel depicts the mean annual maximum daily temperature, by county. The bottom panel depicts the annual number of all crimes per 100,000 persons, by county. All statistics are based on data from 1960-2009.

Source: Matthew Ranson, 'Crime, Weather, and Climate Change', *Journal of Environmental Economics and Management* 67, no. 3 (1 May 2014): 274–302, <https://doi.org/10.1016/j.jeem.2013.11.008>.

The cross-country correlation between temperature and homicide is *negative* across the 37 European countries with populations greater than one million ($r = -0.43$, 95% confidence interval: $-0.66, -0.12$),⁶⁴⁷ i.e. the higher temperature is, the lower is crime.

There is a mild cross-country correlation between homicide and latitude in the Northern Hemisphere, but almost none for the Southern Hemisphere

⁶⁴⁷ Martin Daly and D. B. Krupp, 'The Importance of Being Explicit', *Behavioral and Brain Sciences* 40 (2017).

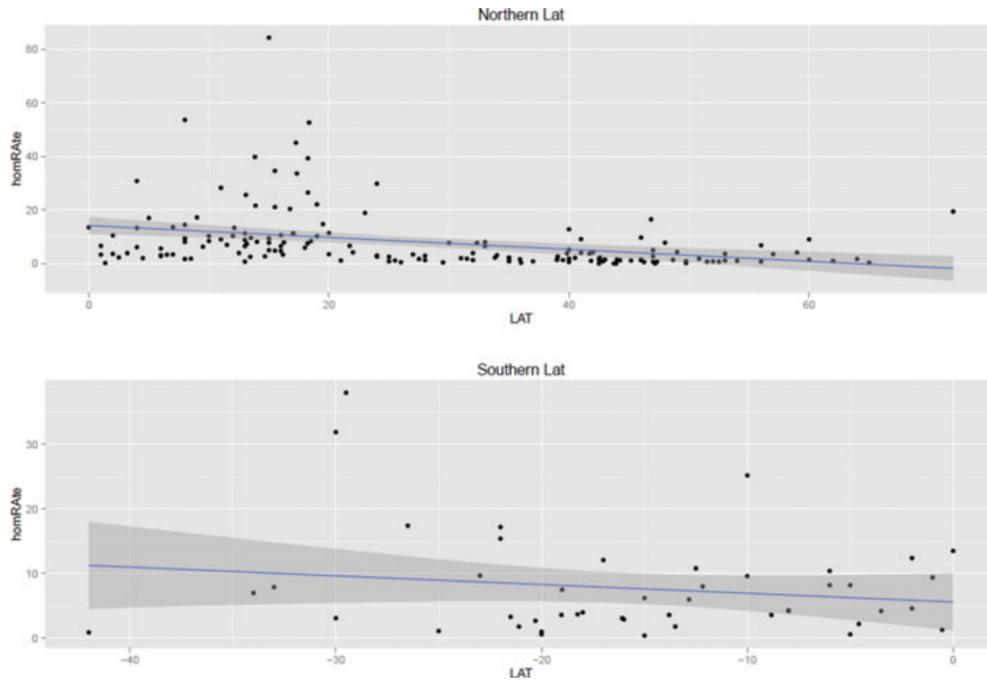


Figure 1. (Fuentes et al.). Homicide rate with respect to latitude for each country.

Source: Agustin Fuentes et al., 'The CLASH Model Lacks Evolutionary and Archeological Support', *Behavioral and Brain Sciences* 40 (2017).

Across Russia, there is a negative correlation between regional violent crime rates and average annual temperature:

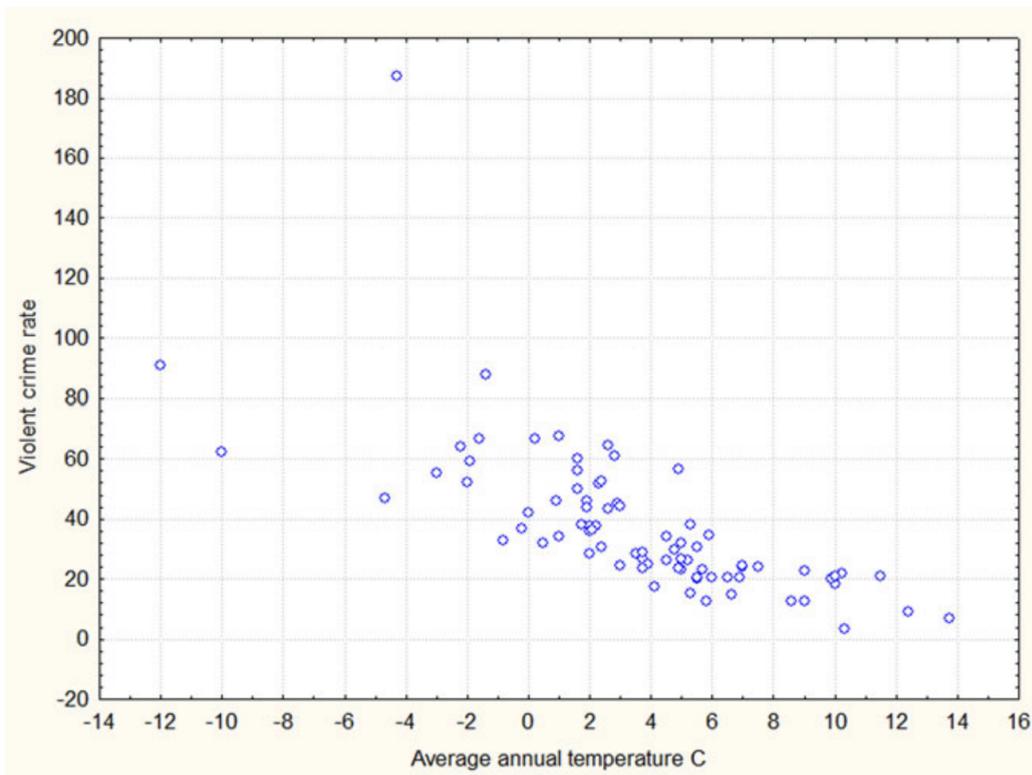


Figure 1. (Prudkov & Rodina). Average annual temperature versus violent crime rate.

Source: Pavel N. Prudkov and Olga N. Rodina, 'Russian Data Refute the CLASH Model', *Behavioral and Brain Sciences* 40 (2017).

Of course, there are lots of confounders for these relationships. But this is indicative evidence that temperature is not a major driver of rates of violence or aggression.

Regarding the issue of confounders, consider this example. Rates of violent crime are higher in Southern US states compared to Northern states, but some scholars attribute this to honour culture, which developed in response to the herding economy of frontier states in the South.⁶⁴⁸ As discussed above, trying to control for confounders in cross-sectional studies is likely to suffer from omitted variable bias and over-controlling.

13.3. Experimental and quasi-experimental evidence

The connection between temperature and interpersonal violence and aggression seems to be much less controversial than the connection between temperature and intergroup conflict, with most research finding a positive relationship.⁶⁴⁹

⁶⁴⁸ Paul AM Van Lange, Maria I. Rinderu, and Brad J. Bushman, 'Aggression and Violence around the World: A Model of CLimate, Aggression, and Self-Control in Humans (CLASH)', *Behavioral and Brain Sciences* 40 (2017): 5.

⁶⁴⁹ "Research has found overwhelming evidence for a positive relationship between violence and temperature." Dennis M. Mares and Kenneth W. Moffett, 'Climate Change and Interpersonal Violence: A "Global" Estimate and Regional Inequities', *Climatic Change* 135, no. 2 (1 March 2016): 297–310, <https://doi.org/10.1007/s10584-015-1566-0>.

There are experimental studies which suggest that higher temperatures increase aggression.

“Other experiments have found that participants are more likely to be hostile, more likely to perceive others as being hostile, and act more aggressively toward another person after being randomly assigned to sit in an uncomfortably hot room instead of a room at a comfortable temperature. One study manipulated temperature in a police training program. They found those police officers randomly assigned to an uncomfortably hot condition were more likely to respond to a burglary scenario by drawing their weapon and by firing it than those who were in a comfortable-temperature condition.”⁶⁵⁰

There is also strong evidence that property and violent crime is highly seasonal, with more crime occurring in the summer.⁶⁵¹

There are also studies exploring natural experiments or quasi-experiments in order to isolate the causal effect of temperature. As discussed above, these studies can correct for the omitted variable and over-controlling problems by holding cultural and political effects constant and exploiting exogenous variation in weather.

“In such time-period studies, we also find a consistent pattern: the hotter the period, the higher the rate of violence. As noted, this can range from a few hours to days, months, or years.”⁶⁵²

“This behavior might contribute to results of natural experiments that use versions of Equation 3, in which it is generally found that crimes between individuals - particularly violent crimes such as assault, murder, rape, and domestic violence - tend to increase at higher temperatures in Australia (Auliciems & DiBartolo 1995), India (Blakeslee & Fishman 2014, Iyer & Topalova 2014), Mexico (Baysan et al. 2014), the Philippines (Wetherley 2014), Tanzania (Burke et al. 2014), and the United States (Anderson et al. 1997, 2000; Cohn & Rotton 1997; Rotton & Cohn 2000; Bushman et al. 2005; Jacob et al. 2007; Card & Dahl 2011; Mares 2013; Ranson 2014).”⁶⁵³

The meta-analysis by Burke et al (2015b) collects together studies of this type and finds that per standard deviation increase in temperature, there is a 2.3% increase in interpersonal

⁶⁵⁰ Andreas Miles-Novelo and Craig A. Anderson, ‘Climate Change and Psychology: Effects of Rapid Global Warming on Violence and Aggression’, *Current Climate Change Reports* 5, no. 1 (2019): 36–46.

⁶⁵¹ “Put simply, crime has a strong seasonality that needs to be properly accounted for in any examination of the relationship between weather and crime” Ryan D. Harp and Kristopher B. Karnauskas, ‘The Influence of Interannual Climate Variability on Regional Violent Crime Rates in the United States’, *GeoHealth* 2, no. 11 (2018): 356–69, <https://doi.org/10.1029/2018GH000152>.

⁶⁵² Andreas Miles-Novelo and Craig A. Anderson, ‘Climate Change and Psychology: Effects of Rapid Global Warming on Violence and Aggression’, *Current Climate Change Reports* 5, no. 1 (2019): 36–46.

⁶⁵³ Marshall Burke, Solomon M. Hsiang, and Edward Miguel, ‘Climate and Conflict’, *Annual Review of Economics* 7, no. 1 (2015): 577–617, <https://doi.org/10.1146/annurev-economics-080614-115430>.

conflict, while a standard deviation change in rainfall causes a 0.59% increase in interpersonal violence.⁶⁵⁴

13.3.1. Methodological issues

Some of the methodological issues with the intergroup conflict studies recur for the interpersonal violence studies.

Weather different to climate

Panel studies test weather variability rather than climate change, and these processes might be importantly different.

Heavy US focus

The studies are quite geographically limited and heavily US focused. Of the 18 interpersonal conflict studies, 9 are based in the US, with the rest covering Australia, India, Mexico and other countries.⁶⁵⁵ The relevance for the rest of the world is therefore not entirely clear.

Implausible precision

The estimates of the mean effect in Burke et al (2015b) seem implausibly precise. They estimate an effect of around 2.3%, with almost no chance that the true effect is <1% or >4%. This level of precision seems implausible given the noisiness of the dependent and independent variable and the many other factors that could affect these variables. I think what is going on here is that they are estimating an effect size on the assumption that their model is correct, but in reality there is much more uncertainty about their model than there is conditional on their model.

Reporting bias

As discussed in the Chapter 12, Burke et al (2015b) find evidence of reporting bias, but do not correct for it. This likely biases the estimate upwards. One can see the publication bias visually in the figure below.

⁶⁵⁴ Marshall Burke, Solomon M. Hsiang, and Edward Miguel, 'Climate and Conflict', *Annual Review of Economics* 7, no. 1 (2015): Table 2, <https://doi.org/10.1146/annurev-economics-080614-115430>.

⁶⁵⁵ Marshall Burke, Solomon M. Hsiang, and Edward Miguel, 'Climate and Conflict', *Annual Review of Economics* 7, no. 1 (2015): Table 1, <https://doi.org/10.1146/annurev-economics-080614-115430>.

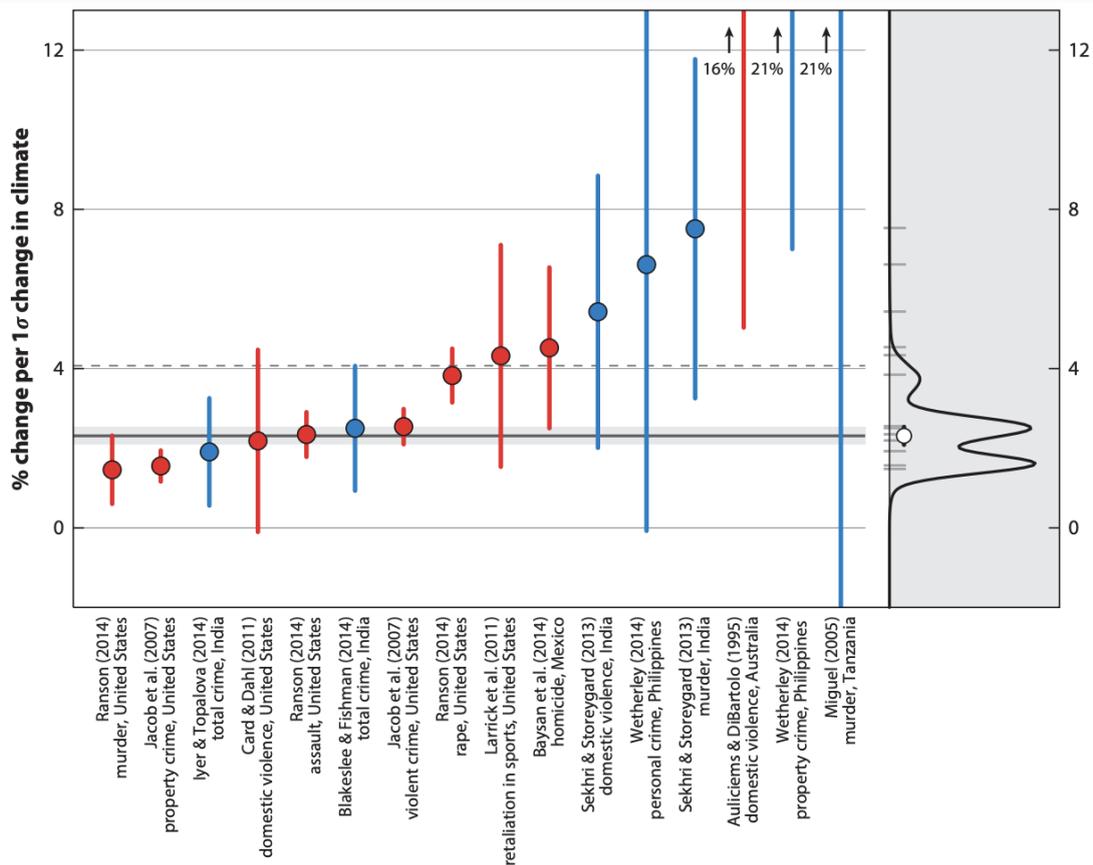


Figure 2

Estimates for the effect of climatic events on the risk of interpersonal violence using authors' preferred specifications. Each marker represents the estimated effect of a 1σ increase in a climate variable, expressed as a percentage change in the conflict variable relative to its mean, and whiskers represent the 95% confidence interval. Colors indicate temperature (*red*) or rainfall loss (*blue*). The dashed line is the median estimate, and the dark gray line is the precision-weighted mean with its 95% confidence interval shown in gray. The panel on the right shows the precision-weighted mean effect (*circle*) and the distribution of study results (*gray ticks*); probability distributions are the posterior for the expected distribution of an additional study (*solid black line*).

Source: Marshall Burke, Solomon M. Hsiang, and Edward Miguel, 'Climate and Conflict', *Annual Review of Economics* 7, no. 1 (2015): <https://doi.org/10.1146/annurev-economics-080614-115430>.

The mean effect is shown by the black horizontal line. Since the mean effect found is so close to zero, one would expect studies to be dotted either side of zero. In fact, all of the studies find a positive effect and the vast majority are statistically significant.

13.4. Mechanisms

The two most popular theoretical explanations for the effect of temperature on interpersonal violence are Routine Activity Theory and General Aggression Theory.⁶⁵⁶

Routine Activity Theory - Increasing temperatures likely increases the frequency of interactions between people, thereby raising the number of violent encounters. For

⁶⁵⁶ Dennis M. Mares and Kenneth W. Moffett, 'Climate Change and Interpersonal Violence: A "Global" Estimate and Regional Inequities', *Climatic Change* 135, no. 2 (1 March 2016): 297–310, <https://doi.org/10.1007/s10584-015-1566-0>.

instance, during pleasant temperatures people exit their homes and go to bars or restaurants where they are more likely to encounter others and drink alcohol.

General Aggression Theory - Explains increases in violence as a result of heightened aggressive emotions brought on by warmer conditions. On this theory, heat raises irritability levels, which leads to increased aggressiveness.

There does not appear to be any consensus in the literature about whether these two theories or any others are true.

Adjudicating between the truth of these two theories is relevant to long-termism because if the General Aggression Theory is true, then the results we see in the quasi-experimental studies reflect an increase in aggression that would be across the entire global population if temperatures were to increase. This is a worse outcome than if the Routine Activities Theory is true.

13.5. Effect size

The literature generally suggests that the effect size of climate, temperature, precipitation and weather on aggression and violence is at most very small.

The cross-sectional studies suggest that temperature is not a major driver of interpersonal violence.

According to Ferguson and Dyke (2012), the effect sizes found in laboratory studies of the General Aggression Theory tend to find effect sizes between $r = 0$ and $r = 0.2$, though they don't provide references, which is frustrating.⁶⁵⁷ This means the General Aggression Theory could explain at most 4% of the variation in aggression. Similarly, in a review of the evidence on the General Aggression Theory, Anderson et al (2000), who are generally supportive of General Aggression Theory, note that laboratory studies have produced inconsistent results: "hot temperatures sometimes increase and sometimes decrease aggressive behavior in laboratory settings".⁶⁵⁸ In a meta-analysis of the laboratory studies, Anderson et al note that across the studies "there was no consistent effect on temperature", though there was a consistent effect if other conditions obtained, such as the contextual presence of negative or neutral affect. Of course, it is natural to be concerned about the methodological quality of these sorts of psychology studies.

Burke et al (2015b) estimate that a one standard deviation change in a climate variable causes a 2.3% increase in interpersonal violence. So, on RCP8.5, following an 8 standard deviation in temperature in 80 years, we would expect an 18% increase in interpersonal violence, other things equal.

⁶⁵⁷ Christopher J. Ferguson and Dominic Dyck, 'Paradigm Change in Aggression Research: The Time Has Come to Retire the General Aggression Model', *Aggression and Violent Behavior* 17, no. 3 (2012): 220–28.

⁶⁵⁸ Craig A. Anderson et al., 'Temperature and Aggression', in *Advances in Experimental Social Psychology*, vol. 32 (Academic Press, 2000), 63–133, [https://doi.org/10.1016/S0065-2601\(00\)80004-0](https://doi.org/10.1016/S0065-2601(00)80004-0).

For comparison, as mentioned above, the US homicide rate dropped by 40% between 1990 and 2000, and many European countries have seen comparable declines between 1990 and 2017, despite warming of about 0.5°C. Other factors seem much more important as drivers of aggression. For example, Reyes (2007) argues that a reduction in childhood lead exposure in the 1970s and 1980s accounted for more than half of the violent crime decline of the 1990s.⁶⁵⁹ Levitt argues that increases in the number of police, the rising prison population, the waning crack epidemic and the legalisation of abortion all explain the decline in US crime in the 1990s.⁶⁶⁰

Overall, the effect of temperature on aggression seems like a negative effect of climate change, but one that is very small and will be swamped by the other determinants of aggression.

⁶⁵⁹ James J. Feigenbaum and Christopher Muller, 'Lead Exposure and Violent Crime in the Early Twentieth Century', *Explorations in Economic History* 62 (1 October 2016): 52, <https://doi.org/10.1016/j.eeh.2016.03.002>.

⁶⁶⁰ Steven D. Levitt, 'Understanding Why Crime Fell in the 1990s: Four Factors That Explain the Decline and Six That Do Not', *Journal of Economic Perspectives* 18, no. 1 (March 2004): 163–90, <https://doi.org/10.1257/089533004773563485>.

Appendix 1. The WCRP estimate of climate sensitivity

Prior selection

What we ought to believe depends on the beliefs we had prior to considering new evidence, and the new evidence. This is the insight of [Bayesian epistemology](#). Many of the models that input into the IPCC's Fifth Assessment Report estimates of climate sensitivity have very heavy right tails, suggesting a non-negligible probability of very high climate sensitivity. This seems to be mainly a product of the fact that these posterior distributions are updated from a uniform prior over climate sensitivity with an arbitrary cut-off at 10°C or 20°C. I checked some of the papers for IPCC models of climate sensitivity that have a long tail and they either: explicitly use a uniform prior which makes a large difference to tail behaviour,⁶⁶¹ or do not say whether or not they use a uniform prior (but I would guess that they do). When this is combined with the likelihood ratio from the data and evidence that we have, we end up with a posterior distribution with heavy right tails.

However, as Annan and Hargreaves (2011) have argued, the use of a uniform prior is unjustified. Firstly, climate scientists use these priors on the assumption that they involve “zero information”, but this is not the case. Secondly and relatedly, the cut-off is arbitrary. Why not have a cut-off at 50°C?

Thirdly, it is not the case that before analysing modern instrumental and paleoclimatic data on the climate, we would rationally believe that a doubling from pre-industrial levels of 280ppm to 560ppm would be equally likely to produce warming of 3°C or 20°C. In 1896, Arrhenius estimated that climate sensitivity was between 5 and 6°C. In 1906, he revised this down to 4°C. Before analysing modern data sets, scientists had settled on a 67% confidence range of between 1.5 and 4.5°C in 1979.⁶⁶² As Annan and Hargreaves note:

“This estimate was produced well in advance of any modern probabilistic analysis of the warming trend and much other observational data, and could barely have been affected by the strong multidecadal trend in global temperature that has emerged since around 1975. Therefore, it could be considered a sensible basis for a credible prior to be updated by recent data.”⁶⁶³

Arguments from physical laws also suggest that extreme values of 10°C or 20°C are extremely unlikely.

⁶⁶¹ Olson Roman et al., “A Climate Sensitivity Estimate Using Bayesian Fusion of Instrumental Observations and an Earth System Model,” *Journal of Geophysical Research: Atmospheres* 117, no. D4 (February 21, 2012), <https://doi.org/10.1029/2011JD016620>; Lorenzo Tomassini et al., “Robust Bayesian Uncertainty Analysis of Climate System Properties Using Markov Chain Monte Carlo Methods,” *Journal of Climate* 20, no. 7 (April 1, 2007): 1239–54, <https://doi.org/10.1175/JCLI4064.1>.

⁶⁶² J. D. Annan and J. C. Hargreaves, “On the Generation and Interpretation of Probabilistic Estimates of Climate Sensitivity,” *Climatic Change* 104, no. 3–4 (February 1, 2011): 429–30, <https://doi.org/10.1007/s10584-009-9715-y>.

⁶⁶³ Annan and Hargreaves, 429–30.

Sherwood et al (2020) corrects for this by testing the exploring the implications of a range of different priors.

Using multiple lines of evidence

Many lines of evidence are relevant to climate sensitivity, including (1) the physics of feedback processes, (2) the historical climate record (since 1800), and (3) the paleoclimate record (millions of years into the past). Because the underlying science spans many disciplines, individual scientists generally only fully understand one or a few of the strands. Consequently, scientists have tended to produce estimates of climate sensitivity using only one strand of evidence. Each approach therefore leaves out potentially important information. Sherwood et al (2020) corrects for this problem by formally taking account of all relevant lines of evidence.⁶⁶⁴

Formally producing a posterior distribution

Our posterior beliefs are what we believe once we have updated our prior beliefs given new evidence. For instance, in advance of considering the historical data, we might believe that climate sensitivity is between 2 and 7°C. Once we learn of new evidence, this might narrow our uncertainty to between 3 and 6°C. The correct way to update our prior beliefs is through formal Bayesian updating.

The last IPCC report produced its estimate of climate sensitivity in an informal way, while trying to account for all relevant lines of evidence.⁶⁶⁵

⁶⁶⁴ “Quantifying ECS is challenging because the available evidence consists of diverse strands, none of which is conclusive by itself. This requires that the strands be combined in some way. Yet, because the underlying science spans many disciplines within the Earth Sciences, individual scientists generally only fully understand one or a few of the strands. Moreover, the interpretation of each strand requires structural assumptions that cannot be proven, and sometimes ECS measures have been estimated from each strand that are not fully equivalent. This complexity and uncertainty thwarts rigorous, definitive calculations and gives expert judgment and assumptions a potentially large role.” p2

⁶⁶⁵ “As a consequence, in this chapter, statements using the calibrated uncertainty language are a result of the expert judgement of the authors, combining assessed literature results with an evaluation of models demonstrated ability (or lack thereof) in simulating the relevant processes (see Chapter 9) and model consensus (or lack thereof) over future projections” IPCC, Climate Change: The Physical Science Basis, 1040.