What We Know and Don't Know about Climate Change, and Implications for Policy

Robert S. Pindyck, *Sloan School of Management, Massachusetts Institute of Technology, and NBER,* United States of America

Executive Summary

There is a lot we know about climate change, but there is also a lot we don't know. Even if we knew how much CO_2 will be emitted over the coming decades, we wouldn't know how much temperatures will rise as a result. And even if we could predict the extent of warming that will occur, we can say very little about its impact. I explain that we face considerable uncertainty over climate change and its impact, why there is so much uncertainty, and why we will continue to face uncertainty in the near future. I also explain the policy implications of climate change uncertainty. First, the uncertainty (particularly over the possibility of a catastrophic climate outcome) creates insurance value, which pushes us to earlier and stronger actions to reduce CO_2 emissions. Second, uncertainty interacts with two kinds of irreversibilities: CO_2 remains in the atmosphere for centuries, making the environmental damage from CO_2 emissions irreversible, pushing us to earlier and stronger actions and reducing CO_2 emissions requires sunk costs, that is, irreversible expenditures, which pushes us away from earlier actions. Both irreversibilities are inherent in climate policy, but the net effect is ambiguous.

I. Introduction

There is a lot we know about climate change, but there is also a lot we don't know. Even if we knew exactly how much carbon dioxide (CO_2) and other greenhouse gases (GHGs) the world will emit over the coming decades, we wouldn't be able to predict with any reasonable precision how much the global mean temperature will rise as a result. Nor would we be able to predict other aspects of climate change, such as rises in sea

levels and increases in the frequency and intensity of storms, hurricanes, and droughts. And even if we were able to predict the extent of climate change that will occur over the coming decades, we can say very little about its likely impact—which in the end is what matters. The fact is that we face considerable uncertainty over climate change, and as we'll see, that uncertainty has crucial implications for policy.

Despite the uncertainty, the debate over climate policy is usually framed in deterministic terms. We start with some scenario regarding GHG emissions, perhaps under "business as usual" or under some emission abatement policy, and then make and discuss projections of temperature change through the end of the century. Sometimes those projections include high, medium, and low alternatives, but without much basis for how and why those alternatives differ as they do. We then talk in broad terms about the likely impacts of those temperature changes—reductions in agricultural output, reduced productivity generally, greater damage from more intense storms and droughts, and perhaps displacements of populations if rising sea levels inundate low-lying areas. We sometimes try to translate those impacts into percentage reductions in gross domestic product (GDP), which is necessary if we want to come up with a number for the social cost of carbon (SCC). We know that those impacts are very difficult—perhaps impossible—to predict because climate change happens slowly, over decades, and we don't know the extent of adaptation that will occur in response.

And despite all the uncertainty, we evaluate climate change policies in terms that suggest a high level of precision is possible. As I have argued elsewhere, this is particularly true when we use complex integrated assessment models (IAMs) to make outcome and impact projections, evaluate alternative policies, and estimate the SCC.[1] But as I will argue, it is the uncertainty over climate change and its impact that is critical to policy formulation, and that should be the focus of analysis and discussion.

To get a sense of why the uncertainties are so important, consider the irreversibilities that are an inherent part of climate policy (and environmental policy more generally). It has been long understood that environmental damage can be irreversible, which can lead to a more "conservationist" policy than would be optimal otherwise. Thanks to Joni Mitchell, even noneconomists know that if we "pave paradise and put up a parking lot," paradise may be gone forever. And because the value of paradise to future generations is uncertain, the benefit from protecting it today should include an option value, which pushes the cost-benefit calculation toward protection. But there is a second kind of irreversibility

that works in the opposite direction: protecting paradise over the years to come imposes sunk costs on society. If paradise includes clean air and water, protecting it could imply sunk cost investments in abatement equipment and an ongoing flow of sunk costs for more expensive production processes. This kind of irreversibility would lead to policies that are less "conservationist" than they would be otherwise.

Which of these two irreversibilities applies to climate policy? Both. Given that they work in opposite directions, which one is more important? We don't know.[2] Because CO_2 can remain in the atmosphere for centuries, and ecosystem destruction from climate change can be permanent, there is clearly an argument for taking early action. But the costs of reducing CO_2 emissions are largely sunk, which implies an argument for waiting.[3] Which type of irreversibility will dominate depends in part on the nature and extent of the uncertainties involved, and will be explored in this paper.

There is another reason why the uncertainties over climate change are so important, and it has to do with "tail risk." If climate change turns out to be moderate, and its impact turns out to be moderate, we may not have too much to worry about. But what if climate change and its impact turn out to be catastrophic—the far right tail of the outcome distribution. It is that possibility, even if the probability is low, that might drive us to quickly adopt a stringent emission abatement policy. In effect, by reducing emissions now we would be buying insurance. But how much of a premium should we be willing to pay for such insurance? The answer depends in part on society's degree of risk aversion, which is complex and hard to evaluate. As I will show, however, the risk premium could be considerable.

This paper has two main parts. First, I lay out what we know, don't know, and sort of know about climate change and discuss why we don't know certain things and the nature of the uncertainties. One of the two more important uncertainties pertains to the extent of warming (and other aspects of climate change) that will occur given current and expected future GHG emissions. The second uncertainty pertains to the economic impact of any climate change that might occur, an impact that depends critically on the possibility of adaptation. Although various estimates are available, we simply don't know how much warmer the world will become by the end of the century under the Paris Agreement, or under any other agreement. Nor do we know how much worse off we will be if the global mean temperature increases by 2°C or even 5°C.

In fact, we may never be able to resolve these uncertainties (at least over the next few decades). It may be that the extent of warming and its impact are not just unknown but also unknowable—what King (2016)

refers to as "radical uncertainty," or extreme Knightian uncertainty.[4] And as King (2016, 131) puts it (in a very different context), "The fundamental point about radical uncertainty is that if we don't know what the future might hold, we don't know, and there is no point pretending otherwise."[5] But even though we may never resolve these uncertainties, we can characterize them and better understand them.

That leads to the second part of this paper, which deals with the implications of uncertainty for climate policy. In a risk-neutral world with no irreversibilities, only the expected values of outcomes should matter, not the degree of uncertainty over those outcomes. But macroeconomic and financial market data suggest that society (or at least the people that make up society) is far from risk-neutral, so that there is likely to be a significant insurance value to reducing GHG emissions now. Likewise, we know that there are two types of irreversibilities at play, which work in opposite directions. In formulating climate policy, what is the insurance value of GHG emission reductions, and what is the net effect of the relevant irreversibilities? This paper addresses those questions.

In the next three sections, I lay out the steps through which emissions of CO_2 (and other GHGs) accumulate in the atmosphere, how increases in the atmospheric CO_2 concentration affect the global mean temperature (and regional temperatures), how temperature increases affect sea levels as well as other aspects of climate, and how changes in climate can in turn have economic and social impacts (i.e., "damages"). I will characterize in general terms the state of our knowledge with respect to each of these steps, that is, the extent of our uncertainty. For two of these steps—how rising GHG concentrations affect climate and how climate change causes damages—the uncertainty is huge.

I will also refine the statement that "the uncertainty is huge." I will try to characterize these uncertainties in terms of probability distributions that have come out of recent studies in climate science and economics. I will address the question of whether those distributions have "fat tails" (and whether that matters). I will also review the evidence on how these uncertainties are changing over time. (As I will explain, between the 2007 and 2014 Intergovernmental Panel on Climate Change [IPCC] reports, uncertainty over how changes in the atmospheric CO_2 concentration affect temperature has actually increased.) This is important because it addresses the value of waiting for new information rather than taking immediate action now.

I will then turn to the implications of uncertainty for policy. First, how does climate change uncertainty interact with the two opposing irreversibilities outlined above? I will address this question using a

simple two-period example. Second, I will explain how climate change uncertainty creates an insurance value of early action. But readers hoping that I can tell them exactly how large that insurance value is will be disappointed. The reason is that there is a catch-22 at work here: the very uncertainties over climate change that create a value of insurance prevent us from determining how large that value is with any precision. On the other hand, we can get a rough sense of how important that insurance value is, and determine whether it is something we should take into account. As we will see, it is indeed something we should take into account.

II.   Some Climate Change Basics

To keep things simple, I will ignore methane and other non-CO_2 GHGs in this paper and focus only on CO_2, which is by far the greatest driver of climate change. Yes, the warming potential of a ton of atmospheric methane is about 25 times the warming potential of a ton of CO_2, but far fewer tons of methane are emitted each year, and methane stays in the atmosphere only for a decade or so, whereas CO_2 stays there for centuries. As a result, methane accounts for less than 10% of the total warming effects of GHG emissions.

It will be useful to go over the basic mechanisms by which CO_2 emissions originate and accumulate in the atmosphere, how increases in the atmospheric CO_2 concentration leads to climate change, how climate change in turn leads to impacts, and how those impacts can be evaluated in economic terms. We also want to know how emissions can be reduced, and at what cost. We could think about this in terms of a projection of climate damages over the coming century under "business as usual," in which nothing is done to reduce emissions, and under alternative emission reduction policies. The steps would be as follows:

1. **GDP Growth:** GHG emissions are generated by economic activity. If all economic activity stopped—no production, no consumption—emissions caused by humans would likewise stop. So the first step in projecting CO_2 emissions is to project GDP growth over the coming century. Not easy! Projecting GDP growth for different countries or regions over the next 5 years is hard enough. (For example, no one anticipated the deep worldwide recession caused by the COVID-19 pandemic.) And now think about projecting GDP growth over the next 50 years. Tough job, and clearly subject to considerable uncertainty.

2. **CO_2 Emissions:** Marching ahead, let's assume we have a reasonable projection of GDP growth (by region) through the end of the century. We would use this information to make projections of future CO_2 emissions under "business as usual," that is, no emission reduction policy, or under one or more abatement scenarios. To do this, we might relate CO_2 emissions to GDP and then use our projections of future GDP. But this is problematic, in part because the relationship between CO_2 emissions and GDP has been changing, and is likely to continue to change in ways that are not entirely predictable. (The impact of the COVID-19 pandemic is an example of how the relationship between CO_2 emissions and GDP can change suddenly and unpredictably.) Note that CO_2 emissions are measured in billions of metric tons, called gigatons (Gt) for short.

3. **Atmospheric CO_2 Concentration:** Suppose we have projections of CO_2 emissions through the end of the century. We could use those projections to project future atmospheric CO_2 concentrations, accounting for past and current emissions as well as future emissions. The key fact is that 1 Gt of CO_2 emitted into the atmosphere increases the CO_2 concentration by 0.128 parts per million (ppm).[6] There is some uncertainty here, because the CO_2 dissipation rate—in the range of 0.0025–0.0050 per year on average—depends in part on the total concentrations of CO_2 in the atmosphere and in the oceans. But relative to other uncertainties, translating emissions to concentrations can be done with reasonable accuracy.

4. **Temperature Change:** Now we come to the hard part. We would like to make projections of the average global mean temperature change likely to result from higher CO_2 concentrations. That means we need a number for *climate sensitivity*—the increase in the global mean temperature that would eventually result from a doubling of the atmospheric CO_2 concentration. OK, so what's that number? Unfortunately, we don't know the true value of climate sensitivity. The "most likely" range (according to the IPCC) is from 1.5°C to 4.5°C, and if we include what the IPCC considers "less likely" but possible values, the range would run from 1.0°C to 6.0°C. Even 1.5°C–4.5°C is a huge range, and it implies a huge range for temperature change. On top of that uncertainty, what is the time lag between an increase in the CO_2 concentration and its impact on temperature? Something like 10–40 years, but again, that is a wide range.

5. **Impact of Climate Change:** But let's march ahead and assume we know how much the temperature will increase during the coming decades

(and how much sea levels will rise, etc.) and try to project the economic impact of such changes in terms of lost GDP and consumption. Now we are in truly uncharted territory. Most IAMs make such projections by including a "damage function" that relates temperature change to lost GDP, but those damage functions are not based on any economic (or other) theory or much in the way of empirical evidence. They are essentially arbitrary functions, made up to describe how GDP goes down when temperature goes up. To make matters worse, "economic impact" should include indirect impacts, such as the social, political, and health impacts of climate change, which might somehow be monetized and added to lost GDP. Here, too, we are in the dark. Basically, we know very little about what the true damage function looks like. The bottom line: projecting the impact of climate change is the most speculative part of the analysis.

6. **Abatement Costs:** To evaluate a candidate climate policy, we must compare the benefits of the policy to its costs. What are the benefits? A reduction in climate-induced damages; for example, a reduction in the loss of GDP that would otherwise result from climate change. But as I just said, projecting the impact of climate change is highly speculative. And what about the costs of a candidate climate policy, that is, the costs of abating GHG emissions by various amounts, both now and throughout the future. A small amount of abatement (say, reducing CO_2 emissions by 5% or 10%) is fairly easy, but a large amount (say, cutting emissions in half) is likely to be quite costly. But how costly? We're not sure, in part because we have had no experience cutting emissions by half or more. Also, we expect that abatement costs will fall over the coming decades, but by how much? Answering that question requires projections of technological change that might reduce future abatement costs, and technological change is hard to predict. Once again, we face considerable uncertainty.

7. **Valuing Current and Future Losses of GDP:** Finally, let's assume that we could somehow determine the annual economic losses (measured in terms of lost GDP) resulting from any particular increase in temperature. Let's also assume that we know the increases in temperature that would result from "business as usual" and under some abatement policy. And suppose we also know the annual cost (again in terms of lost GDP) of that abatement policy. How would we compare the benefits from the policy to its costs? We would need to know the discount rate that would let us compare current losses of GDP (from the cost of abating emissions) with the future gains in GDP from the reduction in damages resulting from the abatement policy.[7] The discount rate (in this

case the social rate of time preference, because it measures how society values a loss of GDP and hence consumption in the future versus today) is critical: a low discount rate (say around 1%) makes it easier to justify an immediate stringent abatement policy; a high rate (say around 5%) does the opposite. So what is the "correct" discount rate? There is no clear number on which economists agree. The US government's Interagency Working Group used three discount rates to estimate the SCC, 2.5%, 3.0%, and 5.0%, although Stern (2015) argues that the "correct" discount rate is about 1.1%. The problem is that 1.1% and 5% will give wildly different estimates of the SCC.

To summarize, there are aspects of climate change—CO_2 emissions and concentrations—where we have a reasonable amount of knowledge and can make reasonable projections. Yes, there is uncertainty, especially when projecting out 50 or more years. But at least we can pinpoint the nature of the uncertainty, and to some extent bound it. And then there are aspects of climate change—changes in temperature, and most notably, the economic impact of those changes—where we know very little. I turn now to a more detailed discussion of what we know and don't know, why we don't know certain things, and the extent of the uncertainty.

III.   What We Know (or Sort of Know)

We understand fairly well some parts of the climate change process. There is uncertainty over the specific numbers, but at least we can estimate those numbers and come up with reasonable bounds.

A.   What Drives CO_2 Emissions?

How much carbon will be burned over the coming decades, and how much CO_2 will be emitted? Putting aside efforts at emissions abatement for now, the answer depends in part on economic activity. As economic activity grows, CO_2 emissions will grow as well. But the answer also depends on the relationship between GDP and CO_2 emissions, and that relationship is neither simple nor fixed. Over the past 50 years or so, the amount of CO_2 emitted per dollar of GDP has declined steadily—in the United States, in Europe, in China, in almost all countries. This ratio, CO_2 emitted per dollar of GDP, is called carbon intensity.

Carbon intensity has been declining for several reasons: (a) The composition of GDP has been changing. Compared with 50 years ago, services have become more important than manufacturing, and services

use less energy and therefore emit less CO_2 than manufacturing. (b) Technological improvements in the way we produce and utilize goods and services have resulted in the use of less energy, and thus lower emissions of CO_2. For example, cars, trucks, and buses are much more fuel efficient than they were 50 years ago, as are home and commercial heating and cooling systems. (c) Energy itself is becoming "greener." Energy generation from renewables (especially wind and solar) has been growing, and the share of energy coming from fossil fuels, especially coal, has been falling.

Carbon intensity and its components can be measured and understood as follows:

1. **Energy Intensity:** The amount of energy consumed per dollar of GDP. We measure energy consumption in quadrillions of BTUs (10^{15} BTUs, denoted as quads), and GDP in billions of US dollars.[8] So the unit of measurement for energy intensity is quad BTUs/$ billion.

2. **Energy Efficiency:** Sometimes referred to instead as CO_2 efficiency, this is the amount of CO_2 emitted from the consumption of 1 quad of energy. If, for example, the energy is generated from wind or solar, little or no CO_2 will be emitted, but a large amount is emitted if the energy is from coal. For energy efficiency, we measure CO_2 emissions in megatons (Mt, millions of metric tons), so the unit of measurement is Mt of CO_2/quad BTUs.

3. **Carbon Intensity:** The amount of CO_2 emitted per $ billion of GDP. Carbon intensity is simply the product of energy intensity and energy efficiency:

$$\text{Carbon Intensity} = \text{Mt } CO_2/\$ \text{ billion}$$
$$= (\text{quads}/\$ \text{ billion}) \times (\text{Mt } CO_2/\text{quad})$$

Decomposing carbon intensity into its two components is useful because the drivers of energy intensity and energy efficiency can be quite different.

What does this decomposition of carbon intensity tell us? It says that if we want to predict CO_2 emissions over the coming decades (with or without some abatement policy), we would have to (1) predict GDP growth, (2) predict changes in energy intensity, and (3) predict changes in energy efficiency. And we'd have to do this for every major country, or at least different regions of the world, because GDP growth, energy intensity, and energy efficiency are likely to evolve very differently in different countries and regions.

What has happened to carbon intensity and its components over the past 40 or 50 years, and what is likely to happen in the future? Briefly:

1. **Energy Intensity:** Figure 1 shows the evolution of energy intensity since 1980 for the world, and for the United States, Europe, India, and China. For the United States and Europe, energy intensity has declined steadily, largely due to gradual changes in the composition of GDP and the ways in which GDP is produced and consumed. Compared with 1980, services are now a larger share of GDP, and the production of services uses less energy than the production of manufactured goods. In addition, we now use less energy to produce and utilize goods and services; cars and trucks have become more fuel efficient, as have household appliances and home and commercial heating and cooling systems. In China, energy intensity has declined sharply, in part because the Chinese GDP was so low in 1980. But there has been little or no decline
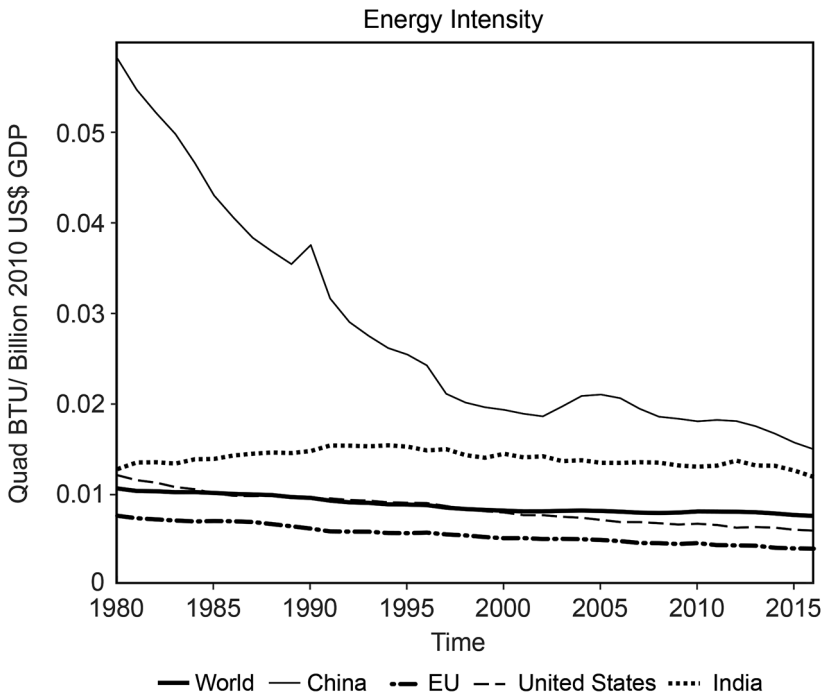

Fig. 1. Energy intensity for the world, and for the United States, European Union (EU), India, and China. Energy intensity is measured in quad (10^{15}) British thermal units (BTUs) per billion 2010 US dollars of gross domestic product (GDP). Sources: World Bank, US Energy Information Agency.

in energy intensity in India and other large developing countries. For the world as a whole, reductions in energy intensity have been quite limited; a decline from 0.011 in 1980 to about 0.0075 quads/$ billion today.

2. **Energy Efficiency:** Even if energy intensity remains constant, we would see a reduction in carbon intensity if we could achieve a significant improvement in energy efficiency. Figure 2 shows the evolution of energy efficiency since 1980 for the world, and for the United States, Europe, India, and China. Both Europe and (to a lesser extent) the United States have had improvements in energy efficiency, in part because energy production is becoming "greener." Energy generation from renewables has been growing, and the share of energy coming from fossil fuels, especially coal, has been falling. But alas, energy efficiency in China and India is now about where it was in 1980—around 70 Mt CO_2/quad BTU in China and around 80 Mt CO_2/quad BTU in India—and well above the levels in the
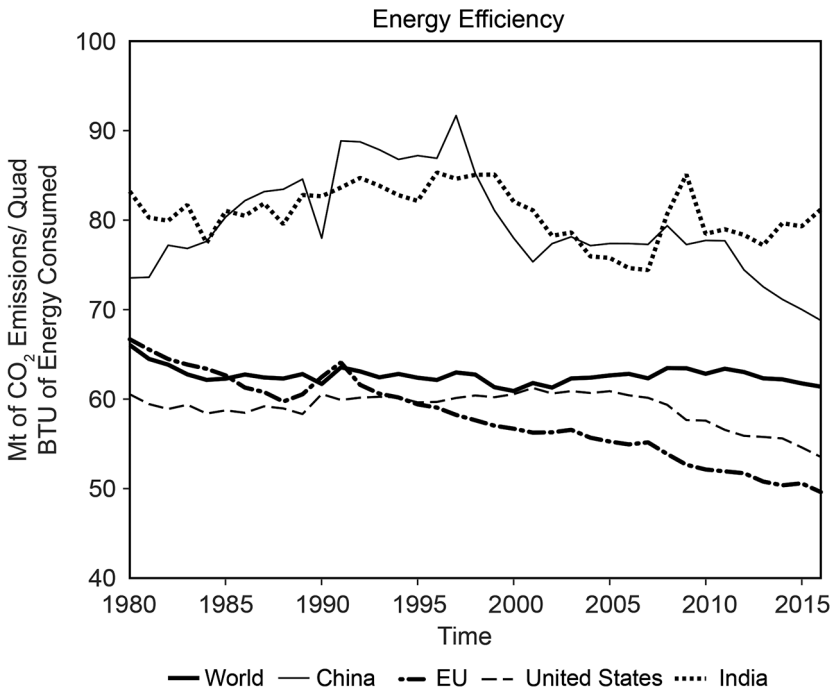

Fig. 2. Energy efficiency for the world, and for the United States, European Union (EU), India, and China. Energy efficiency is measured in megaton (Mt) of CO_2 emissions per quad British thermal unit (BTU) of energy consumed, so a *reduction* in energy efficiency implies an improvement, that is, a reduction in the amount of CO_2 generated from the use of energy. Source: US Energy Information Agency.

United States and Europe. Energy efficiency has followed a similar pattern in other large developing countries. The net result: on a worldwide basis, energy efficiency has remained roughly constant (at about 60 Mt CO_2/ quad BTU).

3. **Carbon Intensity:** What matters in the end is the product of energy intensity and energy efficiency, namely carbon intensity. It has followed a path that is very similar to energy intensity, because energy efficiency hasn't changed much. As illustrated in figure 3, for the world as a whole there has been a gradual decline in carbon intensity from about 0.69 Mt CO_2/$ billion in 1980 to 0.50 Mt CO_2/$ billion in 2000, but after 2000 just a minimal decline, to about 0.46 Mt CO_2/$ billion in 2018.

What does a decline in worldwide carbon intensity from 0.69 Mt CO_2/$ billion in 1980 to about 0.46 Mt CO_2/$ billion in 2018 imply for
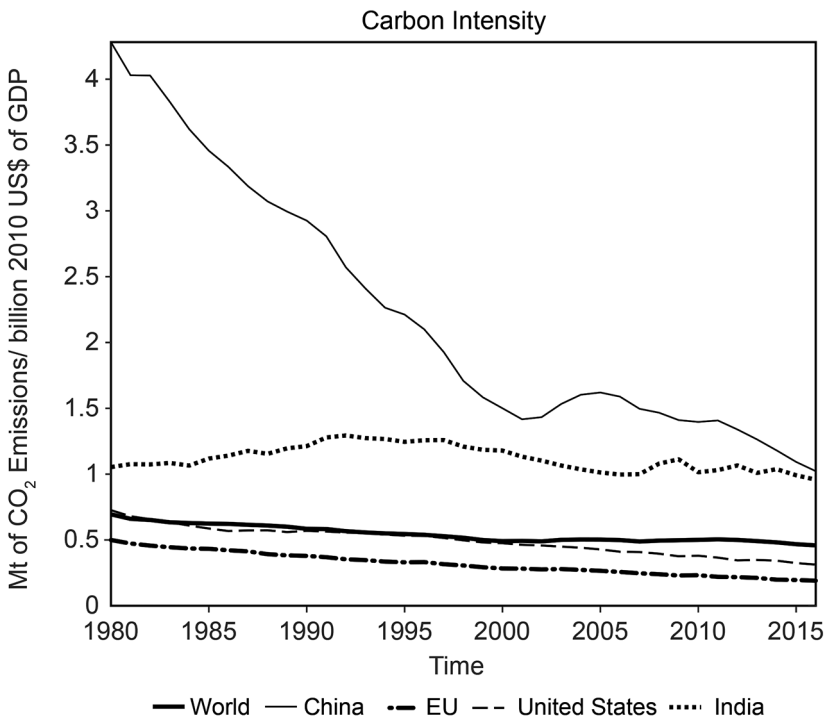

Fig. 3. Carbon intensity for the world, and for the United States, European Union (EU), India, and China. Carbon intensity is the product of energy intensity and energy efficiency, and is measured in megaton (Mt) of CO_2 emissions per billion 2010 US dollars of gross domestic product (GDP).

worldwide CO_2 emissions? If world GDP had remained constant over that time period, CO_2 emissions would have declined by about a third. But (fortunately) world GDP has grown substantially. Measured in 2010 constant US dollars, it nearly tripled, going from about $28 trillion in 1980 to about $80 trillion in 2018. And that's why global CO_2 emissions have increased so much.[9]

What does this tell us about future CO_2 emissions? On one level it paints a rather grim picture. Worldwide, carbon intensity has been declining very slowly, only by about 1% per year, but world GDP has been growing at an average rate of about 3% per year. So there are only two ways that CO_2 emissions can decline in the future: (1) a decline in world GDP or (2) a decline in worldwide carbon intensity. A decline in world GDP is not a happy thought, and we certainly wouldn't want to engineer a global recession as a means of reducing CO_2 emissions. So that leaves us with the second option—a decline in carbon intensity. Do we have any good reasons to expect this to occur? Both energy intensity and energy efficiency can both be affected by government policy. The adoption of strong CO_2 abatement policies seems quite likely in Europe, but less so in the United States, and much less so in key countries such as China, India, Indonesia, and Russia. And the free-riding problem reduces the political feasibility of strong abatement policies in many countries.

Where does this leave us with regard to future CO_2 emissions? If we could predict the growth of GDP around the world and predict the changes in energy intensity and energy efficiency, and thus the changes in carbon intensity, we could come up with at least a rough prediction of future CO_2 emissions. And we would want to make that rough prediction under "business as usual" and under one or more CO_2 abatement policies. Yes, lots of uncertainty, but relatively manageable.

B.   *What Drives the Atmospheric CO_2 Concentration?*

Remember that CO_2 emissions do not directly cause increases in temperature. Instead, warming is caused by increases in the atmospheric CO_2 concentration. Of course increases in the CO_2 concentration are the result of CO_2 emissions, so if we want to make predictions about increases in temperature, we need to determine how any particular path for emissions affects the future path of the CO_2 concentration.

Isn't the current atmospheric CO_2 concentration just the sum of past emissions, minus any dissipation? Roughly, but not precisely. The problem is that some atmospheric CO_2 is absorbed by the oceans, and some

of the CO_2 in the oceans can reenter the atmosphere. How much goes each way? That depends on a variety of factors, including the amounts of CO_2 both in the atmosphere and in the oceans, and the ocean temperature. So even if we had precise projections of CO_2 emissions over the next several decades, our projection of the atmospheric CO_2 concentration would be subject to some uncertainty. Nonetheless, compared with some of the other uncertainties we face, this one is not too bad. Given a predicted path for CO_2 emissions, we can predict the atmospheric CO_2 concentration reasonably well, using the fact that 1 Gt of CO_2 emissions adds 0.128 ppm to the atmospheric CO_2 concentration. Adding up past CO_2 emissions and subtracting dissipation:

$$M_t = (1 - \delta)M_{t-1} + E_t, \tag{1}$$

where E_t is emissions in year t, M_t is the concentration, and δ is the dissipation rate. And what is the correct value for the dissipation rate? Estimates generally range from 0.0025 to 0.0050 per year. Fitting equation (1) to data on CO_2 emissions and the CO_2 concentration yields an estimate of $\delta = .0035$ per year.[10] So once again, although there is some uncertainty, it is relatively manageable.

IV. What We Don't Know

Now we come to the hard part. We would like to make projections of the average global temperature changes likely to result from higher CO_2 concentrations. And then given projections of how much the temperature will increase during the coming decades (and how much sea levels will rise, etc.), what matters is the impact of those changes. If we had reason to believe that higher temperatures and higher sea levels will cause little damage, it would be hard to argue that we should devote resources today on preventative measures. On the other hand, if the likely damages are extreme, then we certainly should act quickly to reduce emissions and prevent climate change. Thus it is important to determine the likely economic impact of warming, rising sea levels, and other measures of climate change in terms of lost GDP and consumption. Unfortunately, when it comes to the impact of climate change we are very limited in what we know, and thus for the most part we can only speculate.

Why is it so difficult to pinpoint climate sensitivity, or at least narrow the range of estimates? Why can't we predict the likely impact of climate change on the economy? I turn now to these questions.

A. Climate Sensitivity

Recall that climate sensitivity is defined as the temperature increase that would eventually result from a doubling of the atmospheric CO_2 concentration. The word "eventually" means after the world's climate system reaches a new equilibrium. It would take a very long time, however, for the climate system to completely reach a new equilibrium, around 300 years or more. However, the climate system will get quite close to equilibrium in a few decades. How many decades depends in part on the size of the increase in the CO_2 concentration—the larger the increase, the longer is the time lag—and even for a given increase, there is some uncertainty over the time lag. But generally 10–40 years is a reasonable range, and 20 years is a commonly used number.[11]

  I said that there is a great deal of uncertainty over the true value of climate sensitivity. Three questions come up. First, just how much uncertainty is there? Second, has research in climate science during the past few decades resulted in more precise estimates of climate sensitivity? In other words, has the uncertainty been reduced, and if so, by how much? And third, why is there so much uncertainty over climate sensitivity? I address each of these questions in turn.

How Much Uncertainty Is There?

Over the past 2 decades, there have been a large number of studies by climate scientists on the magnitude of climate sensitivity. Virtually all of those studies conclude by providing a range of estimates, often in the form of a probability distribution. From the probability distribution we can determine the probability that the true value of climate sensitivity is above or below any particular value, or within any interval; for example, above 4.0°C, or between 2.0°C and 3.0°C. Thus each study gives us an estimate of the nature and extent of uncertainty, according to that study. But there is considerable dispersion across studies, and that dispersion gives us further information regarding the extent of the uncertainty.

  To explore this dispersion, I used information from the roughly 130 studies of equilibrium climate sensitivity assembled by Knutti, Rugenstein, and Hegerl (2017). Most of these studies provide a "best" (most likely) estimate of climate sensitivity, as well as a range of "likely" (i.e., probability greater than 66%) values. Although Knutti et al. (2017) surveyed a few earlier studies, I included only those from 1970 through 2017. I also located and added nine additional studies published in 2017 and 2018.[12]

For each study, I used both the low end of the range of likely values (which I refer to as "minimum estimates") and the high end ("maximum estimates"), as well as the "best" (most likely) estimate. To see how views about climate sensitivity might have changed over time, I divided the studies into two groups based on year of publication: pre-2010 and 2010 onward. Figure 4 shows a histogram with the "best" estimates from these studies.

From the figure, note that the bulk of the studies (115 of the 131) have "best estimates" between 1.5°C and 4.5°C, which is the "most likely" range according to the IPCC. But this is still a wide range, and 16 studies have "best estimates" outside this range (as low as 0.5°C and as high as 8°C). We can also get a sense of how views about climate sensitivity changed by comparing the pre-2010 studies with those published from 2010 onward. Both the mean and standard deviation are somewhat higher for the more recent studies: 2.77 and 1.03, respectively for the pre-2010 studies, and 2.87 and 1.11, respectively, for the later studies.
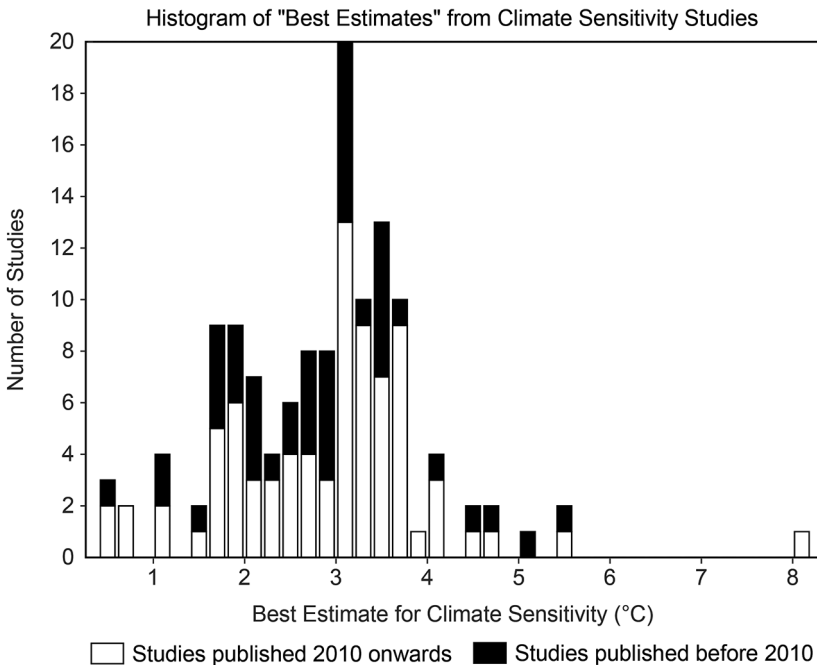

Fig. 4. Histogram of best estimates of climate sensitivity from 131 studies, of which 47 were published prior to 2010 and 84 from 2010 onward. The studies are from Knutti et al. (2017), supplemented by nine additional studies published in 2017 and 2018, and listed in footnote 12.

Figure 5 shows histograms for the low end of the range of likely values reported by these studies ("minimum estimates") and the high end ("maximum estimates"). The bulk of the "minimum estimates" are in the range of 0.5°C–4.0°C, with only three estimates above this range. The bulk of the "maximum estimates" are in the range of 3.0°C–7.0°C, but there are 13 estimates above this range, with seven estimates at 10°C–15°C.

Figure 5 tells us that there is a huge amount of uncertainty over climate sensitivity. If we ignore the outliers and simply consider the bulk of the "minimum" and "maximum" estimates, we get a range of 0.5°C–7.0°C. Remember that this is a range of "likely" (i.e., probability greater than 66%) values, and excludes more extreme values that are unlikely but still possible.

Climate scientists have conducted numerous studies that try to estimate climate sensitivity. The individual studies show large ranges of "likely" values, and that range becomes much greater once we account for the dispersion across the different studies. The bottom line: We are quite certain that climate sensitivity is a positive number, but at this point we simply don't know its actual value. And that's unfortunate, because climate sensitivity is a critical determinant of the temperature increases we can expect over the coming decades.

Has the Uncertainty Been Reduced?

Climate scientists have been busy, publishing hundreds of papers that directly or indirectly relate to climate sensitivity. There is little question that our understanding of the physical mechanisms that underlie climate sensitivity has improved considerably over the past couple of decades. Doesn't this mean that we are now better able to pinpoint the magnitude of climate sensitivity, that is, that our uncertainty over its true value has been reduced?

Unfortunately, the answer is no. In fact, if anything the extent of the uncertainty has grown. This is suggested by the earlier (pre-2010) and later (2010 onward) "best estimates" in the set of studies shown in figure 4; although the distributions are skewed right, the standard deviation is higher for the more recent studies (1.13 vs. 1.03).

The increase in uncertainty has also been demonstrated in a paper by Freeman, Wagner, and Zeckhauser (2015), who compared the survey of climate sensitivity studies in the 2007 IPCC report with the updated survey in the 2014 report. In the 2007 report, the IPCC surveyed 22 peer-reviewed
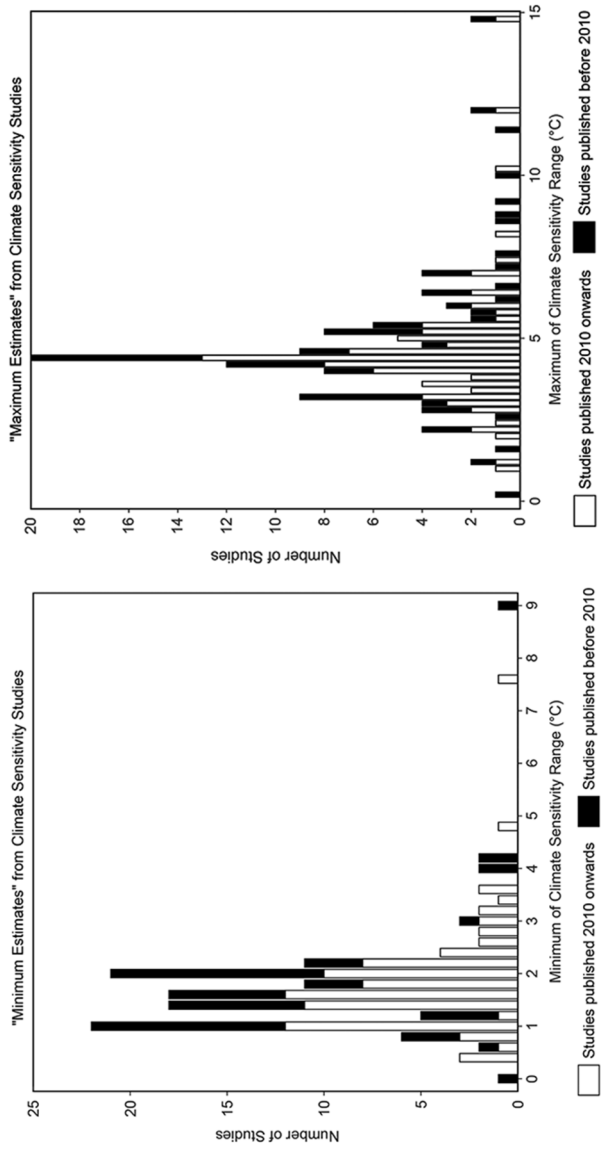
Fig. 5. Histograms of minimum and maximum estimates of climate sensitivity from 143 studies, of which 54 were published prior to 2010 and 89 from 2010 onward. The studies are from Knutti et al. (2017), supplemented by nine additional studies published in 2017 and 2018, and listed in footnote 12.

published studies of climate sensitivity and estimated that the "most likely" range is from 2.0°C to 4.5°C.[13] But then in the 2014 report, that "most likely" range widened, to 1.5°C–4.5°C. Furthermore, the implied standard deviation increased. These results are in part good news, because the bottom of the range became lower (1.5°C instead of 2.0°C). But there is also some bad news, because the estimated uncertainty became greater.

This increase in uncertainty does not mean that climate scientists have not been working diligently, or have otherwise done a bad job. Their work has indeed given us a better understanding of the physical mechanisms through which increases in the atmospheric CO_2 concentration affect temperature. But a better understanding of those physical mechanisms need not mean reduced uncertainty over the magnitude of climate sensitivity. Instead it can simply provide clarity over why there is so much uncertainty.

Why Is There So Much Uncertainty over Climate Sensitivity?

The basic problem is that the magnitude of climate sensitivity is determined by crucial feedback loops, and the parameter values that determine the strength (and even the sign) of those feedback loops are largely unknown, and for the foreseeable future may even be unknowable. This is not a shortcoming of climate science; on the contrary, climate scientists have made enormous progress in understanding the physical mechanisms involved in climate change. But part of that progress is a clearer realization that there are limits (at least currently) to our ability to pin down the strength of the key feedback loops.

The problem is easiest to understand in the context of the simple (but widely cited) climate model of Roe and Baker (2007). It works as follows. Let S_0 represent climate sensitivity in the absence of any feedback effects (i.e., without feedback effects, a doubling of the atmospheric CO_2 concentration would cause an initial temperature increase of $\Delta T_0 = S_0\,°C$). But as Roe and Baker explain, the initial temperature increase ΔT_0 "induces changes in the underlying processes . . . which modify the effective forcing, which, in turn, modifies ΔT." Thus the actual climate sensitivity is given by

$$S = \frac{S_0}{1-f},$$

where f (with $0 \leq f \leq 1$) is the total feedback factor.[14] So if $f = 0.95$, then $S = 20 \times S_0$.

Of course this is an extremely simplified model of the climate system. A more complete and complex model would incorporate several feedback effects; here they are all being rolled into one. Nonetheless, this simple model allows us to address the key problem: Climate sensitivity is very sensitive to the magnitudes of the feedback effects, which in this simplified model comes down to the value of f. But we don't know the value of f. Roe and Baker point out that if we knew the mean and standard deviation of f, denoted by $\bar{f}$ and σ_f respectively, and if σ_f is small, then the standard deviation of S would be proportional to $\sigma_f/(1-\bar{f})^2$. This implies that uncertainty over S is greatly magnified by uncertainty over f, and becomes very large if f is close to 1.

For example, suppose our best estimate of f is 0.95, but we believe that could be off by a factor of 0.03, that is, f could be as low as 0.92 or as high as 0.98. In that case, S could be as low as $(1/.08) \times S_0 = 12.5 \times S_0$ or as high as $(1/.02) \times S_0 = 50 \times S_0$. But $50 \times S_0$ is four times as large as $12.5 \times S_0$, so this seemingly small uncertainty over f creates a huge amount of uncertainty over climate sensitivity.

To illustrate the problem further, Roe and Baker assume that f is normally distributed (with mean $\bar{f}$ and standard deviation σ_f) and derive the resulting distribution for S, climate sensitivity. Given their choice of $\bar{f}$ and σ_f, the resulting median and 95th percentile are close to the corresponding numbers that come from averaging across the standardized distributions summarized by the IPCC.[15]

This Roe-Baker distribution has become well-known and widely used, but it may well understate our uncertainty over climate sensitivity. The reason is that we don't know whether the feedback factor f is in fact normally distributed (and even if it is, we don't know its true mean and standard deviation). Roe and Baker simply assumed a normal distribution. In fact, in an accompanying article in the journal *Science*, Allen and Frame (2007) argued that climate sensitivity is in the realm of the "unknowable," and that the uncertainty will remain for decades to come.

B.   *The Impact of Climate Change*

When assessing climate sensitivity, we at least have scientific results to rely on, and we can argue coherently about the probability distributions that are most consistent with those results. When it comes to predicting the impact of climate change, however, we have much less to go on, and the uncertainty is far greater. In fact, we know very little about the impact

that higher temperatures and rising sea levels would have on the economy and on society more generally.

Why is it so difficult to estimate how climate change will affect the economy? One problem is that we have very few data on which to base empirical work. True, we do have data on temperatures in different locations and in different periods of time, and we can try to relate changes in temperature to changes in GDP and other measures of economic output. In fact there have been some empirical studies that made use of weather data for a large panel of countries over 50 or more years.[16] And there have been many more studies that explore how changes in temperature and rainfall affect agricultural output.[17]

But all of these studies suffer from a fundamental problem: they relate changes in *weather* to changes in GDP or agricultural output, and weather is not the same as climate. The weather in any location (temperature, rainfall, humidity, etc.) changes from week to week and month to month, but the climate—which determines the average temperature and rainfall that we can expect in any particular week or month—changes very slowly (if at all). An unexpectedly hot summer might indeed reduce that year's wheat or corn harvest, but the impact of a gradual change in climate (in which average expected temperatures rise) might have a very different (and probably lower) impact because farmers will shift what and where they plant. Finally, the observed changes in temperature used in these studies are relatively small—not the 4°C or more of warming that many people worry about.

A second problem is that there is little or nothing in the way of economic theory to help us understand the potential impact of higher temperatures. We have some sense of how higher temperatures might affect agriculture, and indeed, most of the empirical work that has been done is focused on agriculture. But we also know that losses of agricultural output in some regions of the world (e.g., near the equator) might be offset by increased output in other regions (e.g., northern Canada and Russia). Furthermore, agriculture is a small fraction of total economic output: 1%–2% of GDP for industrialized countries and 3%–20% of GDP for developing countries. Beyond agriculture, it is difficult to explain, even at a heuristic level, how higher temperatures will affect economic activity.

A third problem is that climate change will occur slowly, which allows for adaptation. This is not to say that adaptation will eliminate the impact of climate change—it will only reduce the impact. But we don't know by how much it will reduce the impact. As a result, adaptation is

another complicating factor that makes it very difficult to estimate the extent of the losses we should expect.

It may be that the relationship between temperature and the economy is not just something we don't know, but something that we cannot know, at least for the time horizon relevant to the design and evaluation of climate policy. As discussed earlier, some researchers have come to the conclusion that climate sensitivity is in this category of the "unknowable," and it may be that the impact of climate change is in that same category. On the other hand, we may start learning more about the impact of climate change, perhaps not in the next few years, but in the next few decades. With more time, and most important with more data related to higher temperatures, it is likely that we will be better able to estimate impacts. For now, however, we need to recognize that our ability to predict the impact of climate change is extremely limited.

C.   A Catastrophic Outcome

It may turn out that over the coming decades climate change and its impact will be mild to moderate. Given all of the uncertainties, this might happen even if little is done to reduce GHG emissions. And if we were certain that this will be case, it would imply that we can relax and stop worrying about climate change.

But we are not certain that the outcome will be so favorable. There is a possibility of an extremely unfavorable outcome, one that we could call catastrophic. Such an outcome would entail a major decline in human welfare from whatever climate change occurs. The IAMs that have been used to make projections have little or nothing to tell us about such outcomes. This is not surprising; the damage functions in these models, which are ad hoc, are calibrated to give small damages for small temperature increases and cannot tell us very much about the kinds of damages we should expect for temperature increases of 5°C or more. And that's unfortunate, because it is the possibility of a catastrophic outcome that really drives the SCC and matters for climate policy.

For climate scientists, a "catastrophic outcome" usually means a high temperature outcome. How high? There is no fixed rule here. Almost all would agree that a 5°C or 6°C increase by 2100 would be in the realm of the catastrophic, and might result if the climate system reaches a "tipping point" as the CO_2 concentration keeps increasing. Putting aside the difficulty of estimating the probability of that outcome, what matters in the end is not the temperature increase itself but rather its impact.

Would that impact be "catastrophic," and might a smaller (and more likely) temperature increase, perhaps 3°C or 4°C, be sufficient to have a catastrophic impact? Again, opinions vary. Some have argued that even a 2°C temperature increase would be catastrophic. For example, Carbon-Brief, an interactive collection of 70 peer-reviewed climate studies that show how different temperatures are projected to affect the world, suggests that 2°C of warming could reduce global GDP by 13%. (The website is https://www.carbonbrief.org.)

Why does the possibility of a catastrophic outcome matter so much for climate policy? Because even if it has a low probability of occurring, the possibility of a severe loss of GDP (broadly interpreted) can justify a large carbon tax (or equivalent emission abatement policy). A mild to moderate outcome, on the other hand, is something to which society can respond, in part through adaptation, at a relatively low cost. This means that to a large extent climate policy has to be based on the (small) likelihood of an extreme outcome.

So how likely is a catastrophic outcome, and how catastrophic might it turn out to be? How high can the atmospheric CO_2 concentration be before the climate system reaches a "tipping point," and temperatures rise rapidly? We don't know. We don't know where a "tipping point," if there is one, might lie, and what the impact of a large temperature increase might be. Furthermore, is difficult to see how answers to these questions will become clear in the next few years, despite all of the ongoing research on climate change. We may know much more in the next 20 years, but in the short term, the likelihood and impact of a catastrophic outcome may simply be in the realm of the "unknowable."

V.   The Policy Implications of Uncertainty

The uncertainties discussed above make the design and analysis of climate policy very different from most other problems in environmental economics. Most environmental problems are amenable to standard cost-benefit analysis. Determining the limits to be placed on sulfur dioxide emissions from coal-burning power plants is a good example. These emissions can harm the health of people living downwind, and they also cause acidification of lakes and rivers, harming fish, and other wildlife. We would like to limit these emissions, but doing so is costly because it would raise the price of the electricity produced by the power plant. On the other hand, the benefit of reducing emissions is a reduction in the health problems that they cause, as well as less damage to lakes and rivers.

So how should we decide the extent to which power plant emissions should be reduced? The standard way that economists approach the problem is to compare the cost of any particular emission reduction to the resulting benefit and consider reducing emissions further if the cost is less than the benefit. Of course there will be uncertainties over the costs and benefits of any candidate policy, but the characteristics and extent of those uncertainties will usually be well understood and comparable in nature to the uncertainties involved in other public and private policy or investment decisions. Economists might argue about the details of the analysis, but at a basic level, we're in well-charted territory and we think we know what we're doing. If we come to the conclusion that a policy to reduce sulfur dioxide emissions by some amount is warranted, that conclusion will be seen—at least by most economists—as defensible and reasonable.

But this is not the case with climate change. Climate policy is controversial in part because the uncertainties are so large and thus complicate the policy arguments. There is considerable disagreement among both climate scientists and economists over the likelihood of alternative climate outcomes, especially catastrophic outcomes. There is also disagreement over the framework that should be used to evaluate the potential benefits from an abatement policy, including the discount rate to be used to put future benefits in present value terms. These disagreements make climate policy much less amenable to standard cost-benefit analysis.

So what should we do? Is there a way to properly account for this uncertainty in our models of climate change? How should we handle the possibility of a catastrophic outcome? And how can we account for the insurance value of early action, and the conflicting irreversibilities inherent in climate policy? I don't pretend to provide complete answers to all of these questions, but in what follows I lay out some ways to think about the problem.

A.   The Value of Climate Insurance

Uncertainty over climate change creates insurance value in two ways, and it is important to keep them clear:

1. First, it occurs through the "damage function," that is, the loss of GDP resulting from any particular temperature increase. Although the impact of any increase in temperature is highly uncertain, it is very likely that the damage function is a convex function of temperature, that is, it becomes increasingly steep as the temperature change becomes larger. Put

another way, going from 3°C of warming to 4°C is likely to cause a much larger reduction in GDP than going from 1°C to 2°C. As the temperature increase becomes larger, adaptation becomes more difficult, so the damage from an additional 1°C of warming becomes larger.

2. The second way that uncertainty creates insurance value is through social risk aversion. Risk aversion refers to a preference for a sure outcome rather than a risky outcome, even if that risky outcome has the same expected value as the sure outcome. We do not know what the "correct" social welfare function is, but we expect it to exhibit at least some risk aversion. This means that society as a whole would pay to avoid the risk of a very bad climate outcome.

The Damage (or Loss) Function

To understand how uncertainty, combined with a convex damage function, creates a value of insurance, we'll use a very simple example. We will consider a single point in the future, say the year 2050, and we will ignore the issue of discounting future costs and benefits. For purposes of this illustrative example, I will assume that the percentage loss of GDP resulting from a temperature increase T is given by

$$L(T) = 1 - \frac{1}{(1 + .01T^2)}.$$ (2)

Equation (2) says that $L(0) = 0$, that is, with no temperature increase, there would be no loss of GDP. It also says that $L(2) = .04$, that is, a 2°C temperature increase would result in a loss of 4% of GDP, $L(4) = 14$, that is, a 4°C temperature increase would result in a loss of 14% of GDP, $L(6) = .26$, that is, a 6°C temperature increase would result in a 26% loss of GDP, and so on. Note that each additional 2°C increase in temperature results in a larger and larger additional loss.

   What does this tell us? First, suppose we know for certain that in 2050 the global mean temperature will have increased by 2°C. And using equation (2), suppose we know that this 2°C temperature increase will cause a 4% drop in GDP, compared with what GDP would be without the higher temperature. Ignoring social risk aversion for now, what percentage of GDP should we be willing to sacrifice to avoid this temperature increase? Up to 4%. Hopefully, we could avoid the temperature increase at a cost that is less than 4% of GDP (perhaps by developing and

making use of new energy-saving technologies). But if we had to, we'd be willing to sacrifice up to 4% of GDP.

Now, suppose there is uncertainty over the temperature increase. We think that the temperature might not increase at all, or that it might increase by 4°C, with each outcome having a 50% probability. The expected value of the temperature increase is (0.5)(0) + (0.5)(4) = 2°C, that is, the same as it was in the first case. But now there is uncertainty. How does this change things?

How bad would a 4°C temperature increase be in terms of its impact on GDP? Would it reduce GDP by 8%, that is, twice the 4% drop we said would occur with a 2°C temperature increase? No, we just saw that the impact would be much larger; the damage caused by higher temperatures will rise more than proportionally. Why? Because 4°C of warming is more likely to cause substantial increases in sea levels (e.g., by melting the Antarctic ice sheets), substantial damage to crops, and so forth. We don't know what the impact would be, but using equation (2) we will assume it causes a 14% drop in GDP. In this case, what percentage of GDP should we be willing to sacrifice to avoid the possibility of a 4°C temperature increase?

To answer this, consider the expected size of the impact on GDP. It is (0.5)(0) + (0.5)(14) = 7% of GDP. That says that we should be willing to sacrifice up to 7% of GDP to avoid the 50% chance of a 4°C temperature increase. (Once again, hopefully we can avoid the temperature increase at a cost that is less than 7% of GDP, but if we had to, we'd be willing to sacrifice up to that amount.)

Let's take this one more step. Suppose there is a 75% probability that there will be no temperature increase, and just a 25% chance of an 8°C temperature increase. And suppose that an 8°C temperature increase would be close to catastrophic and, consistent with equation (2), result in a 40% loss of GDP. The expected value of the temperature increase is still 2°C, but the expected impact of this temperature gamble is now (0.75)(0) + (0.25)(40) = 10% of GDP. That says that we should be willing to sacrifice up to 10% of GDP to avoid a 25% chance of an 8°C temperature increase.

These calculations are summarized in table 1. What's going on here is fairly simple: In terms of its impact on GDP, a 4°C temperature increase is more than twice as harmful as a 2°C temperature increase. So even though there is only a 50% chance of the 4°C increase happening, we would sacrifice a lot to avoid the risk. And an 8°C temperature increase is much more than four times as harmful as a 2°C temperature increase.

Table 1
Possible Temperature Outcomes and Economic Impacts

Maximum T Possible	Probability Max T Occurs	Probability of T = 0	% Loss of GDP if Max T Occurs	Expected Loss of GDP (%)
2°C	1	0	4	4
4°C	.5	.5	14	7
8°C	.25	.75	40	10

Note: Impacts are based on the (hypothetical) loss function $L(T) = 1 - [1/(1 + .01T^2)]$, which gives the percentage loss of GDP resulting from a temperature increase T. Note that in each case the expected temperature change is 2°C. GDP = gross domestic product; Max = maximum.

So we would be willing to pay a lot to avoid a very bad outcome, even if that outcome has only a small chance of occurring. For example, how much would we be willing to pay for the first row of table 1 instead of the third row, that is, for a certain temperature increase of 2°C rather than a 75% chance of no temperature increase and a 25% chance of an 8°C increase? From the last column of the table, we would be willing to give up 10% − 4% = 6 % of GDP. Quite a lot!

This is the essence of insurance: we are willing to pay, sometimes a lot, to avoid a very bad outcome, even if that outcome is very unlikely. So we insure our homes against major damage from fire, storms, or floods, we buy medical insurance to cover the cost of a major hospitalization, and we buy life insurance even if we are healthy and expect to live many more years. This framework suggests that we should be willing to pay a considerable amount for insurance against a very bad (even if unlikely) climate outcome.

The Social Welfare Function

These simple calculations suggest that we should be willing to sacrifice quite a bit of GDP (and hence quite a bit of consumption) to insure against the risk of a very bad climate outcome. But we have understated the value of insurance. We focused on the expected loss of GDP but implicitly assumed that a 10% loss of GDP is exactly twice as bad as a 5% loss. In fact, a 10% loss of GDP might be more than twice as bad as a 5% loss. The reason has to do with how people value more (or less) income and consumption.

Suppose your annual disposable (after-tax) income is $60,000. Suppose this income is increased to $70,000, so you now have an additional

$10,000 to spend on things. That might make you very happy. But now suppose your starting income is $160,000, and we add an extra $10,000, for a total of $170,000. The extra $10,000 will still make you happy, but probably not as much as it would if your starting income was only $60,000. We call this a "declining marginal utility of income"; the value (in terms of the satisfaction it provides) of an additional $10,000 of income is lower the higher your starting income is.

This declining marginal utility of income corresponds to risk aversion. You would probably refuse a lottery in which you had a 50-50 chance of winning $10,000 or losing $10,000. The reason is that (for most people) the value of winning $10,000 is less than the lost value of losing $10,000. How much would you have to be paid to agree to take part in that lottery? Is it $2,000, or $3,000? The higher the amount you'd have to be paid, the greater is your risk aversion.[18] You can think of this amount you'd have to be paid as an insurance premium.

How risk averse is society as a whole? That's a complicated question because society is made up of different people with different attitudes toward risk. Financial market data tell us that investors in the aggregate seem to have substantial risk aversion, but not everyone is an investor, and averting climate change is not the same as investing in the stock market.[19]

So what does this tell us about climate policy? If risk aversion for society as a whole is substantial, that would push us further toward a stringent emissions abatement policy. Apart from that, it shows us why the uncertainties over climate change are so important, and in particular why in this modeling framework society should be willing to sacrifice a substantial amount of GDP to avoid the risk of an extremely bad climate outcome. In effect, by reducing emissions now we would be buying insurance. And the value of that insurance could be considerable.

You might be thinking "Well, this is nice. But exactly how large is the value of climate insurance? To what extent does it push us toward early action, and by how much more should we reduce CO_2 emissions if we want to properly account for the insurance value?" Sorry, but I can't provide those numbers. You may be disappointed with that answer, but remember, we don't know much about the actual loss function (the loss function used to generate table 1 is completely hypothetical), nor do we know the extent of social risk aversion. All we can say at this point is that the value of insurance is likely to be substantial, and will push policy toward earlier and more stringent emission abatement.

B.   *The Effects of Irreversibilities*

Environmental damage can sometimes be irreversible, which can lead to a more "conservationist" policy than would be optimal otherwise. If the value of environmental amenities to future generations is uncertain, the benefit from protecting the environment today should include an option value, which accounts for the possibility that future generations will deeply regret irreversible environmental damage. This option value pushes the cost-benefit calculation toward protection.[20]

But environmental protection requires irreversible expenditures, that is, imposes sunk costs on society. This could include sunk cost investments in abatement equipment and an ongoing flow of sunk costs for alternative and perhaps more expensive production processes. If the future value of the environment is uncertain, this would lead to policies that are less "conservationist" than they would be otherwise. Why? Because future generations might find it less valuable than we currently expect, in which case they will regret the irreversible expenditure that we made on preservation.

Given that these two irreversibilities work in opposite directions, which one is more important? We don't know. Because CO_2 can remain in the atmosphere for centuries, and ecosystem destruction from climate change can be permanent, there is clearly an argument for taking early action. But the costs of reducing CO_2 emissions are largely sunk, which implies an argument for waiting.[21] Which type of irreversibility will dominate depends in part on the nature and extent of the uncertainties involved, as we will see.

Before proceeding, it is important to be clear about the nature of "learning" and its connection to climate change uncertainty. Over the next 2 decades, it is likely that our understanding of climate change and its impact will improve considerably. Although so far our uncertainty over climate sensitivity has not decreased (and as discussed above, has actually increased somewhat), more data combined with advances in climate science are likely to reduce the uncertainty. And more data will likely improve our understanding and ability to predict climate change impacts. But at the end of the 2 decades there will still be a good deal of uncertainty as we look toward the next 2 decades. It's a bit like forecasting the price of oil. We don't know what the price will be 5 years from now, but we will find out when the 5 years are up. Then what? As we look out to the next 5 years, there will again be uncertainty. Nonetheless, the ongoing uncertainty creates option value, in this case pushing us

away from investing in the development of new oil reserves or related projects.[22]

Now let's return to climate change policy. The implications of the two conflicting irreversibilities described above can be understood with a simple numerical example.

A Numerical Example

Suppose we must decide whether to spend money now to reduce CO_2 emissions, and then we will decide again in the future, say 40 years from now. We'll assume that at each time there are only two choices: spend nothing on abatement ($A = 0$) or spend 6% of GDP on abatement ($A = .06$). If today we spend nothing ($A_1 = 0$), there will be 10 units of CO_2 emissions that will accumulate in the atmosphere. So, denoting emissions now by E_1 and the atmospheric concentration by M_1, we will have $E_1 = M_1 = 10$. On the other hand, if we do spend 6% of GDP to abate emissions ($A_1 = .06$), emissions will be reduced by 80%, so that $E_1 = M_1 = 2$. Finally, we will assume that CO_2 emissions are partly irreversible: 50% of the today's emissions will dissipate over the next 40 years, so if we emit 10 units of CO_2 today, only 5 units will remain.

To keep this simple, we will also assume that today's emissions cause no damage to the economy now; any damage will occur only in the future. Also, right now we don't know how much damage atmospheric CO_2 will cause in the future: there is a 50% chance that atmospheric CO_2 will cause no damage (the "good" outcome) and a 50% chance it will cause significant damage (the "bad" outcome). Of course 40 years from now there will still be uncertainty over climate change impacts another 40 years out—there will always be uncertainty about future events and impacts. But for purposes of this very simple example, we will only be concerned with decisions now and 40 years from now. The abatement and outcome possibilities are summarized in table 2 and also illustrated in figure 6.

Suppose there is no abatement now ($A_1 = 0$), so 10 units of CO_2 are emitted. How much abatement would we want in the future? The answer depends on the economic impact, which by then we will know. If the impact is zero (the "good" outcome), then there is no reason to abate, so we will have $A_2 = 0$. (This outcome is not shown in the table.) But if the "bad" outcome occurs (an 8% loss of GDP), we will want to abate emissions, that is, set $A_2 = .06$. As table 2 shows, with the "bad" outcome and $A_2 = 0$, the loss of GDP will be 31%, but with $A_2 = .06$,

Table 2

Example Illustrating the Trade-off between Immediate Emissions Abatement versus Waiting for Information about Impact of Warming

% GDP for Abatement, A_1	$M_1 = E_1$	% GDP for Abatement, A_2	$M_2 = (1 - \delta) M_1 + E_2$	"Bad Outcome" Loss of GDP
$A_1 = 0$	10	$A_2 = 0$	5 + 10 = 15	31%
$A_1 = 0$	10	$A_2 = .06$	5 + 2 = 7	**17%**
		Expected Loss if $A_1 = 0$: (.5)(.17) + (.5)(.06) = 11.5%		
$A_1 = .06$	2	$A_2 = 0$	1 + 10 = 11	25%
$A_1 = .06$	2	$A_2 = .06$	1 + 2 = 3	**8%**
		Expected Loss if $A_1 = .06$: (.5)(.08) + (.5)(.06) = 7%		

Note: There are two periods, "now" and, say, 40 years from now. A_1 is expenditure on abatement now, as percentage of GDP, and A_2 is expenditure 40 years from now. We denote emissions by E and the amount in the atmosphere by M. If $A_1 = 0$, there will be 10 units of emissions (E_1), which will accumulate in the atmosphere (so $M_1 = 10$), but half will dissipate over the next 40 years ($\delta = .5$). If $A_1 = .06$ (6% of GDP is spent on abatement), emissions will be reduced by 80%, so that $E_1 = M_1 = 2$. Damages occur in 40 years, and depend only on the CO_2 in the atmosphere at that time, $M_2 = (1 - \delta)M_1 + E_2$. With equal probability the impact could be "good," in which case there is no loss of GDP, or "bad," in which case the loss of GDP is $1 - [1/(1 + .03M)]$, and is shown in the last column. Whatever the value of A_1, if the impact turns out to be "bad," it is best to abate, that is, set $A_2 = .06$, so that the loss of GDP is only 17% (and shown in bold). Also shown is the expected loss of GDP if $A_1 = 0$ (11.5%) and if $A_1 = .06$ (7%). Because the difference (11.5 − 7 = 4.5) is less than the 6% cost of abatement, it is better not to abate now, but instead wait and abate in the future only if we learn the impact is "bad." GDP = gross domestic product.

the loss of GDP will only be 17%. Abatement will cost 6% of GDP, but we will save (31 − 17) = 14 % of GDP, so the investment in abatement is clearly warranted, and the 17% is shown in bold.

Why not set $A_1 = .06$ at the outset, before we learn whether the impact will be "bad" or "good"? Because spending 6% of GDP on abatement is an irreversible expenditure, which we will regret if it turns out the impact is "good." But to see whether the potential regret is large enough, we have to see what happens if we do set $A_1 = .06$ at the outset. As table 2 shows, with $A_1 = .06$, only 2 units of CO_2 will be emitted, and of those 2 units, only 1 will remain after 40 years. If we then learn that the impact is "good," there will be no reasons to abate, so we will set $A_2 = 0$. But if the impact is "bad," it will be best to abate, so we will set $A_2 = .06$.

Now let's come back to the initial decision regarding A_1. What is the expected loss of GDP if we set $A_1 = 0$? As shown in table 2, there is a 50% chance that the impact will turn out to be "bad," in which case we will set $A_2 = .06$ (which costs 6% of GDP) and lose 17% of GDP. So the expected loss if $A_1 = 0$ is (0.5)(.17) + (0.5)(.06) = 11:5%. Also shown is
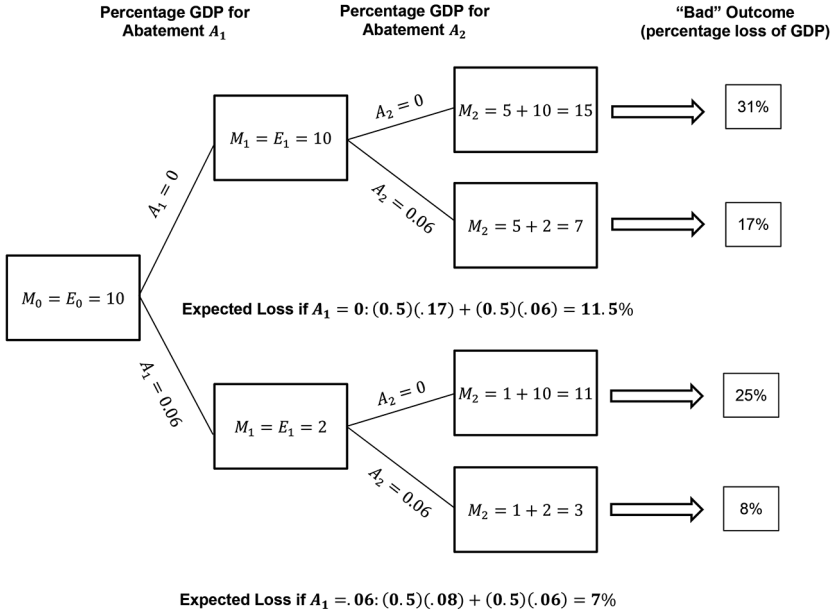
Percentage GDP for Abatement A_1

Percentage GDP for Abatement A_2

"Bad" Outcome (percentage loss of GDP)

$A_1 = 0$

$M_1 = E_1 = 10$

$A_2 = 0$

$M_2 = 5 + 10 = 15$

31%

$A_2 = 0.06$

$M_2 = 5 + 2 = 7$

17%

$M_0 = E_0 = 10$

Expected Loss if $A_1 = 0$: $(0.5)(.17) + (0.5)(.06) = 11.5\%$

$A_1 = 0.06$

$M_1 = E_1 = 2$

$A_2 = 0$

$M_2 = 1 + 10 = 11$

25%

$A_2 = 0.06$

$M_2 = 1 + 2 = 3$

8%

Expected Loss if $A_1 = .06$: $(0.5)(.08) + (0.5)(.06) = 7\%$

Fig. 6. Trade-off between immediate emissions abatement versus waiting. This figure provides another way of looking at the information in table 2.

the expected loss of GDP if A_1 = .06, which turns out to be 7%. Because the difference (11.5 − 7 = 4:5%) is less than the 6% cost of abatement, it is better not to abate now but instead to wait and abate in the future only if we learn the impact is "bad."

To summarize, we have assumed that CO_2 emissions are only partly irreversible, that is, 50% of today's emissions will dissipate over the next 40 years. The cost of abatement (6% of GDP), however, is completely irreversible; it is a sunk cost that can never be recovered. In this case, given that the impact of CO_2 is uncertain and will only be known in the future, it is better to wait, rather than spend 6% of GDP now on abatement. In this case the abatement cost irreversibility outweighs the environmental irreversibility.

Revising the Example

But now let's change one of the key assumptions. This time we will assume that there is no dissipation of CO_2 once it enters the atmosphere. This means setting $\delta = 0$ so that $M_2 = M_1 + E_2$. The results are shown in table 3 and also illustrated in figure 7.

Table 3

Modified Example Illustrating the Trade-off between Immediate Emissions Abatement versus Waiting for Information

% GDP for Abatement, A_1	$M_1 = E_1$	% GDP for Abatement, A_2	$M_2 = (1 - \delta)M_1 + E_2$	"Bad Outcome" Loss of GDP
$A_1 = 0$	10	$A_2 = 0$	$10 + 10 = 20$	37.5%
$A_1 = 0$	10	$A_2 = .06$	$10 + 2 = 12$	**26.5%**
		Expected Loss if $A_1 = 0$: $(.5)(.265) + (.5)(.06) = 16\%$		
$A_1 = .06$	2	$A_2 = 0$	$2 + 10 = 12$	26.5%
$A_1 = .06$	2	$A_2 = .06$	$2 + 2 = 4$	**11%**
		Expected Loss if $A_1 = .06$: $(.5)(.11) + (.5)(.06) = 8.5\%$		

Note: Everything here is the same as in table 2, except the dissipation rate, δ, is 0. Whatever the value of A_1, if in the future the impact turns out to be "bad," it is best to abate, that is, set $A_2 = .06$, yielding the GDP loss shown in bold. Also shown is the expected loss of GDP if $A_1 = 0$ (16%) and if $A_1 = .06$ (8.5%). Now the difference ($16 - 8.5 = 7.5\%$) is greater than the 6% cost of abatement, so it is optimal to abate immediately. Because emissions are now completely irreversible, we are pushed toward early action. The sunk (irreversible) cost of abatement remains, pushing us toward waiting, but now the environmental irreversibility dominates. GDP = gross domestic product.
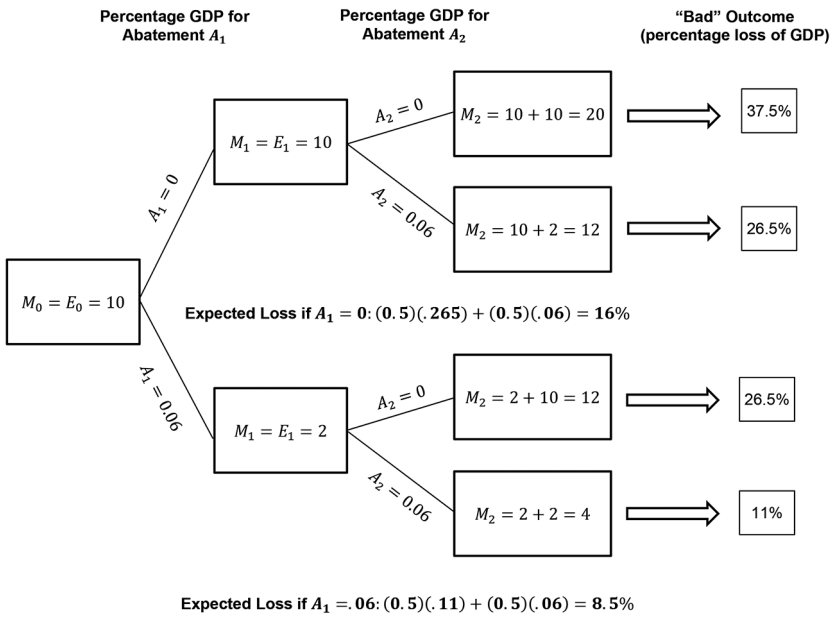

Fig. 7. Modified example of trade-off between immediate emissions abatement versus waiting. This figure provides another way of looking at the information in table 3.

Because we have now assumed that any CO_2 emitted into the atmosphere stays there forever, the loss of GDP under the "bad" outcome will be greater, whatever the abatement policy happens to be. (Compare the last column of table 3 with the last column of table 2.) As in the previous example, whatever the value of A_1, if in the future the impact turns out to be "bad," it is best to abate, that is, set $A_2 = .06$.

What is the optimal abatement policy today? As before, we find out by calculating the expected loss of GDP if we set $A_1 = 0$, and the expected loss if we set $A_1 = .06$. If $A_1 = 0$ the expected loss of GDP is 16%, and if $A_1 = .06$ the expected loss is 8.5%. Now the difference $(16 - 8.5 = 7.5 \%)$ is greater than the 6% cost of abatement, so it is optimal to abate immediately. The uncertainty is the same as before, but because emissions are now completely irreversible (there is no dissipation), we are pushed toward early action. The sunk (irreversible) cost of abatement remains, pushing us toward waiting, but now the effect of the environmental irreversibility dominates.

To further illustrate, these effects of irreversibility, let's make another modification to the numbers in table 3. As in the table, we will assume that there is no dissipation, that is, whatever CO_2 is emitted in the beginning will remain in the atmosphere over the 40 years. However, we will make one simple change: we now assume that positive abatement requires an expenditure of 8% of GDP rather than 6%. In other words, in table 3 and figure 7, replace $A_1 = .06$ with $A_1 = .08$ and $A_2 = .06$ with $A_2 = .08$. Now we can once again calculate the expected loss if $A_1 = 0$ and the expected loss if $A_1 = .08$. Doing so, we will find that the expected loss is 17% if $A_1 = 0$ and 9% if $A_1 = .08$. The difference, $17 - 9 = 8 \%$, is just equal to the 8% cost of abatement. In this case the effects of the two irreversibilities just balance out, so we would be indifferent between abating now and not abating now.

Emissions Abatement: Hold Back or Accelerate?

These numerical examples were simply designed to illustrate the opposing effects of the two irreversibilities that are an inherent aspect of climate policy. But now you might be thinking that the examples are interesting, but what do they tell us about the real world? Which of these two irreversibilities is more important when it comes to actual climate policy? Should we hold back on emissions abatement because of the sunk cost, or should we accelerate abatement because of the irreversible environmental damage caused by emissions? And by how much should we

hold back or accelerate? Sorry, but I can't answer these questions. Why not? Because we simply don't know enough about the climate system and about the impact of varying amounts of climate change.

In the numerical examples I assumed that CO_2 emissions could be reduced by 80% at a cost of 6% (or 8%) of GDP. But we don't actually know how much it would cost (in terms of a percentage of GDP) to reduce emissions by 80%. What we do know is that the cost would be sunk, that is, irreversible, which would lead us to hold back, and the greater the cost the more we would want to hold back. And although we know that CO_2 can remain in the atmosphere for centuries, we don't know what effect it would have on temperature, or what the impact of a higher temperature would be on GDP and other measures of social welfare. (I simply assumed a relationship between the amount of CO_2 in the atmosphere and the percentage reduction in GDP.) These uncertainties, combined with the near-permanence of atmospheric CO_2, would lead us to accelerate abatement.

The balance between these two irreversibilities is also affected by the degree of social risk aversion for the economy as a whole. The sunk cost of abatement can be estimated, at least roughly. But the effect of CO_2 emissions on temperature and the impact of temperature on GDP are highly uncertain. Coming back to our discussion of climate change insurance, these uncertainties would amplify the effect of the environmental irreversibility, and thereby push us toward accelerated abatement.

VI. Conclusions

Unfortunately, many of the books, articles, and press reports that we read make it seem that we know a lot more about climate change and its impact than is actually the case. Likewise, commentators and politicians often make statements of the sort that if we don't take immediate action and sharply reduce CO_2 emissions, the following things will happen, as though we actually know what will happen. Rarely do we read or hear that those things *might happen*; instead we're told they *will happen*.

This shouldn't come as a surprise. We humans prefer certainty to uncertainty, we feel uncomfortable when we don't know what lies ahead, and many people have trouble understanding concepts involving probabilities. Most people prefer to hear or read statements of the sort "By 2050 the temperature will rise by X°C, sea levels will rise by Y meters, and as a result GDP will fall by Z%," as opposed to "There is a 10% chance that temperature will rise by X°C." Many people ignore the fact,

or find it hard to accept, that even if we could accurately predict future GHG emissions, we don't know—and at this point can't know—by how much the temperature or sea levels will rise. And even if we could accurately predict how much the temperature and sea levels will rise, we don't know what the impact would be on GDP or other measures of economic and social welfare. The simple fact is that the "climate outcome," by which I mean the extent of climate change and its impact on the economy and society more generally, is far more uncertain than most people think. This is reflected in the wide variation in expert opinion, as I have shown here in the context of climate sensitivity and I have shown in Pindyck (2019) in the broader context of the SCC.

Our uncertainty over climate change and its impact has important implications for policy. Some would argue that, with so much uncertainty, we should wait and see what happens rather than try to sharply reduce emissions right away. After all, if we don't know how much the climate will change, and we don't know what the impact of climate change will be, why take costly actions now? There is something to that argument, because those costly actions are largely irreversible. But there is another irreversibility that works in the opposite direction, and that is the environmental damage itself; CO_2 emissions remain in the atmosphere for centuries. Which of these two irreversibilities dominates? Unfortunately, we just don't know enough about the climate system to say. (I provided some illustrative numerical examples, but they are just examples.)

There is another reason why uncertainty need not lead us to delay action: with uncertainty, especially the kind of uncertainty we face in the climate sphere, we need insurance. The kind of uncertainty I am talking about is the possibility of a catastrophic outcome, that is, tail risk. I have explained that "climate insurance" is valuable for two reasons. First, although we can't specify the damage function in any detail, we do know that the incremental damage (in terms of lost GDP) from an extra 1°C of warming increases sharply with the total amount of warming. That was the basis for the simple examples in tables 2 and 3. Second, most people exhibit substantial risk aversion, so it is reasonable to think that the social welfare function (representing society as a whole) should also exhibit risk aversion.

So what, exactly, is the value of "climate insurance?" I can't say, because there is a catch-22 at work here: the very uncertainties over climate change that create a value of insurance prevent us from determining exactly how large that value is. This may disappoint some readers, who perhaps were hoping that I would state just how much CO_2 emissions

should be reduced. On the other hand, the simple numerical examples we explored suggest that the insurance value is likely to be large. And what is most important, the very fact that there is an insurance value is a reason why the correct policy response to uncertainty is not to sit back and wait to see what happens.

Endnotes

1. For a discussion of the flaws in IAMs that make them unsuitable for policy analysis, see Pindyck (2013a, 2013b, 2017). The US government's Interagency Working Group used three IAMs to estimate the SCC; see Interagency Working Group on Social Cost of Carbon (2013), and for a discussion of the Working Group's methodology and the models it used, see Greenstone, Kopits, and Wolverton (2013). For a different point of view on the value of IAMs, see Nordhaus (2014) and Weyant (2017).

2. A number of studies have explored this question in a theoretical setting; see, e.g., Kolstad (1996), Ulph and Ulph (1997), and Pindyck (2000). These studies illustrate the fundamental problem, but they don't tell us how to formulate climate policy.

3. There are other arguments for waiting or starting slowly: technological change may reduce abatement costs in the future, and the fact that the "unpolluted" atmosphere is an exhaustible resource implies that the SCC should rise over time (as the atmospheric CO_2 concentration rises).

4. For explanations of why "radical uncertainty" is likely to apply to climate change, see, e.g., Allen and Frame (2007) and Roe and Baker (2007). Heal and Millner (2014) also describe some of the uncertainties inherent in climate change and their implications for policy.

5. I would argue that the IAMs and related models used for policy analysis pretend otherwise, insofar as their projections understate the extent of uncertainty.

6. In 2018, global CO_2 emissions were about 36 Gt, so that year's emissions increased the atmospheric CO_2 concentration by about $(36)(0.128) = 4.61$ ppm.

7. We might also want to specify a social welfare function, i.e., the loss of social utility resulting from a loss of GDP (and hence from a loss of consumption). If GDP and consumption are very high, the loss of utility resulting from a 5% loss of GDP would be smaller than if GDP and consumption were low.

8. One BTU (British thermal unit) is the amount of heat energy required to raise the temperature of 1 lb of water by 1°F. In the metric system, the unit of energy is the calorie, which is the amount of heat required to raise the temperature of 1 g of water by 1°C. One BTU is approximately 252 cal.

9. The connection between GDP growth and climate change is also discussed by Stock (2019), who describes statistical approaches to estimating the relationship.

10. For example, emissions in 1961 were 9 Gt, which added $(9)(0.128) = 1.15$ ppm of CO_2 to the 315 ppm already in the atmosphere. Dissipation in 1961 was $(.0035)(315) = 1.10$ ppm, so the net increase was 0.05 ppm, making the 1961 concentration $315 + 0.05 = 315.05$.

11. Climate scientists often distinguish between "equilibrium climate sensitivity," which is climate sensitivity as I have described it above, and "transient climate response,"

which is the response of global mean temperature to a gradual (1% per year) increase in the CO_2 concentration. See National Academy of Sciences (2017, 88–95) for a discussion. I will simply use the term "climate sensitivity," and treat the time lag (10–40 years) as the time it takes for the climate system to get close to equilibrium.

12. All of the studies that Knutti et al. (2017) examined are listed in their paper. The nine additional studies that I added are Brown and Caldeira (2017), Krissansen-Totton and Catling (2017), Andrews et al. (2018), Cox, Huntingford, and Williamson (2018), Dessler and Forster (2018), Lewis and Curry (2018), Lohmann and Neubauer (2018), Skeie et al. (2018), and Keery, Holden, and Edwards (2018).

13. Intergovernmental Panel on Climate Change (IPCC 2007) also provides a detailed and readable overview of the physical mechanisms involved in climate change, and the state of our knowledge regarding those mechanisms. Each of the individual studies included a probability distribution for climate sensitivity, and by putting the distributions in a standardized form, the IPCC created a graph that showed all of the distributions in a summary form. This is updated in IPCC (2014).

14. In the notation of Roe and Baker (2007), λ_0 is climate sensitivity without feedback effects, and λ is climate sensitivity accounting for feedback effects.

15. Adding a displacement parameter θ, the Roe-Baker distribution is given by:

$$ g(S; \bar{f}, \sigma_f, \theta) = \frac{1}{\sigma_f \sqrt{2\pi}z^2} \exp\left[-\frac{1}{2} \left(\frac{1 - \bar{f} - 1/z}{\sigma_f} \right)^2 \right], $$

where $z = S + \theta$. Fitting to the distributions summarized by the IPCC, the parameter values are $\bar{f} = 0.797$, $\sigma_f = .0441$, and $\theta = 2.13$. This distribution is fat-tailed, i.e., declines to zero more slowly than exponentially. Weitzman (2009, 2014) has shown that parameter uncertainty can lead to a fat-tailed distribution for climate sensitivity, and that this implies a relatively high probability of a catastrophic outcome, which in turn suggests that the value of abatement is high. Pindyck (2011) shows that a fat-tailed distribution by itself need not imply a high value of abatement.

16. For example, Dell, Jones, and Olken (2012) found that the impact of higher temperatures is largely on the growth rate of GDP, as opposed to its level, and is mostly significant in poor countries. See Dell et al. (2014) for an overview of this line of research.

17. For overviews, see Auffhammer et al. (2013) and Blanc and Schlenker (2017).

18. A utility function translates income, or wealth or consumption, into units of well-being (or satisfaction). For a textbook explanation, see Pindyck and Rubinfeld (2018). A commonly used utility function is

$$ u(y) = \frac{1}{1 - \eta} y^{1-\eta}, $$

where y is income and η is called the coefficient of relative risk aversion. In this case marginal utility, i.e., the benefit of an additional dollar of income, is $du/dy = y^{-\eta}$. Marginal utility declines with the level of income, and the larger η is the faster it declines. Thus the larger is η, the greater is the insurance premium you would require to take part in a lottery for which there is a 50–50 chance of winning or losing $10,000.

19. Based on financial market data, and data on consumption and savings, the coefficient of relative risk aversion for society as a whole seems to be in the range of 2–5, which is substantial.

20. One of the earliest studies to analyze this implication of irreversible environmental damage is Arrow and Fisher (1974).

21. There are other arguments for waiting or starting slowly: technological change may reduce abatement costs in the future, and the fact that the "unpolluted" atmosphere is an exhaustible resource implies that the SCC should rise over time (as the atmospheric CO_2 concentration rises).

22. For a textbook treatment of option value, including its application to the development of oil reserves, see Dixit and Pindyck (1994). For an application to environmental policy, see Pindyck (2000).

References

Allen, Myles R., and David J. Frame. 2007. "Call Off the Quest." *Science* 318:582–83.

Andrews, Timothy, Jonathan M. Gregory, David Paynter, Levi G. Silvers, Chen Zhou, Thorsten Mauritsen, Mark J. Webb, Kyle C. Armour, Piers M. Forster, and Holly Titchner. 2018. "Accounting for Changing Temperature Patterns Increases Historical Estimates of Climate Sensitivity." *Geophysical Research Letters* 45 (16): 8490–99.

Arrow, Kenneth J., and Anthony C. Fisher. 1974. "Environmental Preservation, Uncertainty, and Irreversibility." *Quarterly Journal of Economics* 88 (2): 312–19.

Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy* 7:181–98.

Blanc, Elodie, and Wolfram Schlenker. 2017. "The Use of Panel Models in Assessments of Climate Impacts on Agriculture." *Review of Environmental Economics and Policy* 11 (2): 258–79.

Brown, Patrick T., and Ken Caldeira. 2017. "Greater Future Global Warming Inferred from Earth's Recent Energy Budget." *Nature* 552 (7683): 45.

Cox, Peter M., Chris Huntingford, and Mark S. Williamson. 2018. "Emergent Constraint on Equilibrium Climate Sensitivity from Global Temperature Variability." *Nature* 553 (7688): 319.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics* 4:66–95.

———. 2014. "What Do We Learn from the Weather? The New Climate-Economy Literature." *Journal of Economic Literature* 52 (3): 740–98.

Dessler, A. E., and P. M. Forster. 2018. "An Estimate of Equilibrium Climate Sensitivity from Interannual Variability." *Journal of Geophysical Research: Atmospheres* 123 (16): 8634–45.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

Freeman, Mark C., Gernot Wagner, and Richard Zeckhauser. 2015. "Climate Sensitivity Uncertainty: When Is Good News Bad?" *Philosophical Transactions A* 373:1–15.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton. 2013. "Developing a Social Cost of Carbon for Use in U.S. Regulatory Analysis: A Methodology and Interpretation." *Review of Environmental Economics and Policy* 7 (1): 23–46.

Heal, Geoffrey, and Antony Millner. 2014. "Uncertainty and Decision Making in Climate Change Economics." *Review of Environmental Economics and Policy* 8 (1): 120–37.

Interagency Working Group on Social Cost of Carbon. 2013. *Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis*. Washington, DC: US Government.

IPCC (Intergovernmental Panel on Climate Change). 2007. *Climate Change 2007*. Cambridge: Cambridge University Press.

———. 2014. *Climate Change 2014*. Cambridge: Cambridge University Press.

Keery, John S., Philip B. Holden, and Neil R. Edwards. 2018. "Sensitivity of the Eocene Climate to CO_2 and Orbital Variability." *Climate of the Past* 14 (2): 215–38.

King, Mervyn. 2016. *The End of Alchemy: Money, Banking, and the Future of the Global Economy*. New York: W. W. Norton.

Knutti, Reto, Maria A. A. Rugenstein, and Gabriele C. Hegerl. 2017. "Beyond Equilibrium Climate Sensitivity." *Nature Geoscience* 10:727–44.

Kolstad, Charles D. 1996. "Fundamental Irreversibilities in Stock Externalities." *Journal of Public Economics* 60:221–33.

Krissansen-Totton, Joshua, and David C. Catling. 2017. "Constraining Climate Sensitivity and Continental versus Seafloor Weathering Using an Inverse Geological Carbon Cycle Model." *Nature Communications* 8:15423.

Lewis, Nicholas, and Judith Curry. 2018. "The Impact of Recent Forcing and Ocean Heat Uptake Data on Estimates of Climate Sensitivity." *Journal of Climate* 31 (15): 6051–71.

Lohmann, Ulrike, and David Neubauer. 2018. "The Importance of Mixed-Phase and Ice Clouds for Climate Sensitivity in the Global Aerosol-Climate Model ECHAM6-HAM2." *Atmospheric Chemistry and Physics* 18 (12): 8807–28.

National Academy of Sciences. 2017. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*. Washington, DC: National Academies Press.

Nordhaus, William. 2014. "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches." *Journal of the Association of Environmental and Resource Economists* 1 (1/2): 273–312.

Pindyck, Robert S. 2000. "Irreversibilities and the Timing of Environmental Policy." *Resource and Energy Economics* 22:233–59.

———. 2011. "Fat Tails, Thin Tails, and Climate Change Policy." *Review of Environmental Economics and Policy* 5 (2): 258–74.

———. 2013a. "Climate Change Policy: What Do the Models Tell Us?" *Journal of Economic Literature* 51 (3): 860–72.

———. 2013b. "The Climate Policy Dilemma." *Review of Environmental Economics and Policy* 7 (2): 219–37.

———. 2017. "The Use and Misuse of Models for Climate Policy." *Review of Environmental Economics and Policy* 11 (1): 100–114.

———. 2019. "The Social Cost of Carbon Revisited." *Journal of Environmental Economics and Management* 94:140–60.

Pindyck, Robert S., and Daniel L. Rubinfeld. 2018. *Microeconomics*. 9th ed. New York: Pearson.

Roe, Gerard H., and Marcia B. Baker. 2007. "Why Is Climate Sensitivity So Unpredictable?" *Science* 318:629–32.

Skeie, Ragnhild Bieltvedt, Terje Koren Berntsen, Magne Tommy Aldrin, Marit Holden, and Gunnar Myhre. 2018. "Climate Sensitivity Estimates—Sensitivity to Radiative Forcing Time Series and Observational Data." *Earth System Dynamics* 9 (2): 879–94.

Stern, Nicholas. 2015. *Why Are We Waiting? The Logic, Urgency, and Promise of Tackling Climate Change*. Cambridge, MA: MIT Press.

Stock, James H. 2019. "Climate Change, Climate Policy, and Economic Growth." In *NBER Macroeconomics Annual*, ed. Martin S. Eichenbaum, Erik Hurst, and Jonathan A. Parker. Chicago: University of Chicago Press.

Ulph, Alistair, and David Ulph. 1997. "Global Warming, Irreversibility and Learning." *Economic Journal* 107 (442): 636–50.

Weitzman, Martin L. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91:1–19.

———. 2014. "Fat Tails and the Social Cost of Carbon." *American Economic Review* 104 (5): 544–46.

Weyant, John. 2017. "Some Contributions of Integrated Assessment Models of Global Climate Change." *Review of Environmental Economics and Policy* 11 (1): 115–37.