

CLIMATE SHOCK

THE ECONOMIC CONSEQUENCES
OF A HOTTER PLANET

Gernot Wagner
Martin L. Weitzman

PRINCETON UNIVERSITY PRESS
PRINCETON AND OXFORD

Copyright © 2015 by Princeton University Press

Published by Princeton University Press, 41 William Street, Princeton, New Jersey 08540

In the United Kingdom: Princeton University Press, 6 Oxford Street, Woodstock, Oxfordshire OX20 1TW
press.princeton.edu

All Rights Reserved

Library of Congress Cataloging-in-Publication Data

Wagner, Gernot, 1980–

Climate shock : the economic consequences of a hotter planet / Gernot Wagner and Martin L. Weitzman.

pages cm

Includes bibliographical references and index.

ISBN 978-0-691-15947-8 (hardcover : alk. paper) 1. Climatic changes—Economic aspects. I. Weitzman, Martin L., 1942– II. Title.

QC903.W34 2015

363.73874—dc23 2014028540

British Library Cataloging-in-Publication Data is available

This book has been composed in Sabon Next LT Pro, DINPro

Printed on acid-free paper. ∞

Printed in the United States of America

1 3 5 7 9 10 8 6 4 2

Contents



Preface: Pop Quiz	ix
Chapter 1. 911	1
Chapter 2. 411	30
Chapter 3. Fat Tails	48
Chapter 4. Willful Blindness	80
Chapter 5. Bailing Out the Planet	92
Chapter 6. 007	116
Chapter 7. What You Can Do	128
Epilogue: A Different Kind of Optimism	148
<i>Acknowledgments</i>	153
<i>Notes</i>	155
<i>Bibliography</i>	207
<i>Index</i>	243

Preface



Pop Quiz

TWO QUICK QUESTIONS:

- Do you think climate change is an urgent problem?
- Do you think getting the world off fossil fuels is difficult?

If you answered “Yes” to both of these questions, welcome. You’ll nod along, on occasion even cheer, while reading this book. You’ll feel reaffirmed.

You are also in the minority. The vast majority of people answer “Yes” to one or the other question, but not both.

If you answered “Yes” *only* to the first question, you probably think of yourself as a committed environmentalist. You may think climate change is *the* issue facing society. It’s bad. It’s worse than most of us think. It’s hitting home already, and it will strike us with full force. We should be pulling out all the stops: solar panels, bike lanes, the whole lot.

You’re right, in part. Climate change is an urgent problem. But you’re fooling yourself if you think getting off fossil fuels will be simple. It will be one of the most difficult challenges modern civilization has ever faced, and it will require the most sustained, well-managed, globally cooperative effort the human species has ever mounted.

If you answered “Yes” *only* to the second question, chances are you don’t think climate change is the defining problem of our generation. That doesn’t necessarily mean you’re a “skeptic” or “denier” of the underlying scientific

evidence; you may still think global warming is worthy of our attention. But realism dictates that we can't stop life as we know it to mitigate a problem that'll take decades or centuries to show its full force. Look, some people are suffering right now because of *lack* of energy. And whatever the United States, Europe, or other high emitters do to rein in their energy consumption will be nullified by China, India, and the rest catching up with the rich world's standard of living. You know there are trade-offs. You also know that solar panels and bike lanes alone won't do.

You, too, are right, but none of that makes climate change any less of a problem. The long lead time for solutions and the complex global web of players are precisely why we must act decisively, today.



If you are an economist, chances are you answered “Yes” to the second question. Standard economic treatments all but prescribe the stance of the “realist.” After all, economists live and breathe trade-offs. Your love for your children may go beyond anything in this world, but as economists we are obligated to say that, strictly speaking, it's not infinite. As a parent, you may invest enormous sums of money and time into your children, but you, too, face trade-offs: between doing your day job and reading bedtime stories, between indulging now and teaching for later.

Trade-offs are particularly relevant on an average, national, or global level. And they are perhaps nowhere more apparent on the planetary scale than in the case of climate change. It's the ultimate battle of growth versus the environment. Stronger climate policy now implies higher, immediate economic costs. Coal-fired power plants will become obsolete sooner or won't be built in the first place.

That comes with costs, for coal plant owners and electricity consumers alike. The big trade-off question then is how these costs compare to the benefits of action, both because of lower carbon pollution and because of economic returns from investing in cleaner, leaner technologies today.

Economists often cast themselves as the rational arbiters in the middle of the debate. Our air is worse now than it was during the Stone Age, but life expectancy is a lot higher, too. Sea levels are rising, threatening hundreds of millions of lives and livelihoods, but societies have moved cities before. Getting off fossil fuels will be tough, but human ingenuity—technological change—will surely save the day once again. Life will be different, but who's to say it will be worse? Markets have given us longer lives and untold riches. Let properly guided market forces do their magic.

There's a lot to be said for that logic. But the operative words are "properly guided." What precisely are the costs of unabated climate change? What's known, what's unknown, what's unknowable? And where does what we don't know lead us?

That last question is *the* key one: Most everything we know tells us climate change is bad. Most everything we don't know tells us it's probably much worse.

"Bad" or "worse" doesn't mean hopeless. In fact, almost every prediction in this book is prefaced by a version of the words *unless we act*. We don't venture predictions only to see them become true. We talk about where unfettered economic forces may lead in order to guide them in a more productive, better direction. And guide we can. In many ways, putting a proper price on carbon isn't a question of *if*, it's a question of *when*.

CHAPTER 3



Fat Tails

IN 1995, the Intergovernmental Panel on Climate Change (IPCC) declared it was “more likely than not” the case that global warming was caused by human activity. By 2001, it was “likely.” By 2007, it was “very likely.” By 2013, it was “extremely likely.” There’s only one step left in official IPCC lingo: “virtually certain.” The big question is how certain the world needs to be to act in a way that is commensurate with the magnitude of the challenge.

An equally important question is whether all this talk of certainty is conveying what it ought to convey. The increasing likelihood of anthropogenic climate change has three sides to it. Only one of them is good.

The first piece of bad news is that we humans are, in fact, increasing global temperatures and sea levels alike. It would have been cause for celebration if, say, the 2013 report decided that science had gotten it wrong all along. Imagine the *New York Times* headline: “IPCC says decade ‘without warming’ here to stay.” Alas, no such luck. Modern atmospheric science once again confirmed the basic ideas of high school chemistry and physics, going back to the 1800s: more carbon dioxide in the atmosphere traps more heat.

The good news then, in some twisted philosophical sense, is the confirmation of bad news. Climate science has progressed over the past couple decades to the point where it is now able to make the definitive statement that global warming is *extremely likely* caused by human activity. We

know enough to act. Ignoring that reality, by now, would amount to willful blindness.

But there's an additional piece of bad news: the false sense of security conveyed by all this talk about certainty. At least by one important measure, we don't appear to be closer to understanding how much our actions will warm the planet than we were in the 1970s, at the dawn of modern climate science and long before the first IPCC report. Worse still is that what we have learned since then points to the fact that what happens at the very extremes—the tails of the distribution—may dwarf all else.

SENSITIVE CLIMATES

In 1896—eight decades before Wally Broecker coined the term “global warming,” and long before anyone knew what a climate model was—Swedish scientist Svante Arrhenius calculated the effect of doubling carbon dioxide levels in the atmosphere on temperatures. Arrhenius came up with a range of 5 to 6°C (9–11°F). That effect—what happens to global average surface temperatures as carbon dioxide in the atmosphere doubles—has since become known as “climate sensitivity” and has turned into an iconic yardstick.

Climate sensitivity itself is already a compromise, a way of making an incredibly complex topic slightly more tractable. The parameter does have a few things going for it. For one, the starting level of carbon in the atmosphere doesn't matter, at least not by much. One of the few rather well-established facts is that eventual global average temperatures scale linearly with percentage changes in underlying carbon dioxide concentrations. The first 1 percent increase in carbon in the atmosphere has a similar impact

as the 100th. Any doubling of concentrations, from anywhere within a reasonable range, leads to roughly the same eventual increase in global temperatures. The definition of climate sensitivity plays off that fact.

A doubling of preindustrial carbon dioxide levels of 280 parts per million (ppm) seems all but inevitable. The world has just passed carbon dioxide concentrations of 400 ppm, and levels are still rising at 2 ppm per year. Counting other greenhouse gases, the International Energy Agency (IEA) estimates that the world will end up somewhere around 700 ppm by 2100—two-and-a-half times preindustrial levels—unless major emitters take drastic additional steps.

Luckily, Arrhenius's climate sensitivity range of 5–6°C (9–11°F) has proven to be too pessimistic. In 1979, a National Academy of Sciences Ad Hoc Study Group on Carbon Dioxide and Climate concluded that the best estimate of climate sensitivity was 3°C (5.4°F), give or take 1.5°C (2.7°F).

“Conclude” may be a bit strong a term to use here. The process is commonly retold thus, not without admiration for academic genius at work: Jule Charney, the study's lead author, looked at two prominent estimates at the time—2°C (3.6°F) on one end and 4°C (7.2°F) on the other—averaged them to get 3°C (5.4°F), and added half a degree centigrade on either end to round out the range because, well, uncertainty.

Thirty-five years of ever more sophisticated global climate modeling later, our confidence in the range has increased, but what's now called the “likely” range of 1.5 to 4.5 °C (2.7 to 8°F) still stands. That should be a tipoff right there that something rather strange is happening. There's something stranger still.

A PLANETARY CRAPSHOOT

The IPCC defines “likely” events to have at least a 66 percent chance of occurring. That still tells us nothing about whether things may turn out all right—with climate sensitivity closer to 1.5°C (2.7°F)—or not all right at all—closer to 4.5°C (8°F). Taking the IPCC probability descriptions literally, the chance of being outside that range would be up to 34 percent. There’s no precise verdict as to where these 34 percent go, though there’s clearly more room above 4.5°C (8°F) than below 1.5°C (2.7°F). See figure 3.1.

For any numbers below 1.5°C (2.7°F), we could rightfully celebrate—ideally with a bottle of Champagne, flown in from France for the occasion, and letting out an extra puff of carbon dioxide when opened. Though that’s unlikely. And not even these low climate sensitivity realizations of 1.5°C (2.7°F) provide a guarantee that climate change won’t be bad. In fact, quite the opposite: At 700 ppm, final temperatures would still rise to higher than where they were over three million years ago. Think




Figure 3.1 Eventual global average surface warming due to a doubling of carbon dioxide (climate sensitivity)

back to camels in Canada, who happily marched in what is now frozen tundra at temperatures of 2–3.5°C (3.6–6.3°F) above preindustrial levels. And we'd be at 2°C (3.6°F) with a climate sensitivity of 1.5°C (2.7°F), which is at the lower edge of the likely range.

All that makes our inability to exclude climate sensitivities above 4.5°C (8°F) all the more significant. Any probability of climate sensitivity that high should make for (heat-induced) shudders. The most important question then is: how fast does the chance of hitting any of these higher climate sensitivity figures go to zero as the upper bound of climate sensitivity increases? One could imagine an extreme scenario in which the chance that climate sensitivity is above 4.5°C is greater than 10 percent, but if the chance of being above 4.6°C were zero, we could exclude any even higher numbers. If only the planet were that lucky. It's *extremely unlikely*—in the English rather than the strict IPCC sense of that term—that the probabilities of higher climate sensitivities would drop off that quickly.

It's much more likely that the chance of hitting higher temperatures tapers off at an uncomfortably slow pace, before hitting something close enough to zero to provide a reasonable level of comfort that even more extreme numbers won't materialize. That scenario is closer to what statisticians describe as a "fat tail." The probability of 4.6°C is smaller than for 4.5°C, though not by much.

The all-important question, then: how likely is a potentially catastrophic realization of climate sensitivity? The IPCC says it's "very unlikely" that climate sensitivity is above 6°C (11°F). That's comforting but for its definition of just what "very unlikely" means: a chance of anywhere between 0 and 10 percent. And that range is still only the

likelihood that *climate sensitivity* is above 6°C (11°F), not actual temperature rise.



Let’s jump right to the conclusion. Take the latest consensus verdict at face value and assume a “likely” range for climate sensitivity of between 1.5 and 4.5°C (2.7 and 8°F). Equally important, stick to the IPCC definition of “likely” and assume it means a chance of greater than 66 percent, but less than 90 percent. (The latter would be “very likely.”) And take the IEA’s interpretation of current government policy commitments at face value. Here’s what you get: about a 10 percent chance of eventual temperatures exceeding 6°C (11°F), unless the world acts much more decisively than it has.

Figure 3.2 and table 3.1 are the culmination of parsing umpteen scientific papers and countless hours spent fretting over how to get it just so. Rows 1 and 2 in the table




Figure 3.2 Eventual global average surface warming based on passing 700 ppm CO₂e

Table 3.1. Chance of eventual warming of $>6^{\circ}\text{C}$ (11°F) rises rapidly with increasing CO_2e concentrations

<i>CO₂e</i> concentration (ppm)	400	450	500	550	600	650	700	750	800
Median temperature increase	1.3°C (2.3°F)	1.8°C (3.2°F)	2.2°C (4.0°F)	2.5°C (4.5°F)	2.7°C (4.9°F)	3.2°C (5.8°F)	3.4°C (6.1°F)	3.7°C (6.7°F)	3.9°C (7.0°F)
Chance of $>6^{\circ}\text{C}$ (11°F)	0.04%	0.3%	1.2%	3%	5%	8%	11%	14%	17%

represent the move from carbon dioxide–equivalent (CO_2e) concentrations in the atmosphere to ultimate temperature increases. Row 3 shows the corresponding chance of exceeding final average temperature increases of 6°C (11°F). Whenever we had to make a judgment call of where to go next, we tried to take the more conservative turn, which may well underplay some of the true uncertainties involved.

The scariest bit is just how fast the chance of eventual temperatures exceeding 6°C (11°F) goes up. Compare changes in the median temperature increase with the chance of passing 6°C (11°F). Going from 400 to 450 ppm, the difference between 1.3°C (2.3°F) and 1.8°C (3.2°F) for the most likely temperature increase, may not be all that much. There may be some, potentially irreversible, tipping points along the way, but ultimately it's only half a degree centigrade (less than a degree Fahrenheit), or an increase of a bit more than a third. At the same time, the chance of exceeding 6°C (11°F), the last row, just jumped from 0.04 percent to 0.3 percent, almost tenfold. All that's just for moving from 400 to 450 ppm, while the world has already passed 400 ppm for carbon dioxide alone and 440 to 480

ppm for carbon dioxide–equivalent concentrations! A further jump to 500 ppm increases that chance of catastrophe to 1.2 percent. By the time concentrations reach 700 ppm—where the IEA projects the world will end up by 2100 even if all governments keep all their current promises—the chance of eventually exceeding 6°C (11°F) rises to about 10 percent. That looks like the manifestation of a fat tail, if there ever was one (even though strictly speaking we don’t even assume that property in our calculations; our tail is “heavy,” not quite “fat” in statistical terms).

At 700 ppm, the median temperature increase would be 3.4°C (6.1°F). This alone would be a profound, earth-as-we-know-it-altering change. Polar regions would likely warm by at least twice that global average, with everything that entails. The costs would be staggering and should have prompted the world’s leaders to head off such a possibility long ago. Yet those costs are still nothing compared to what would happen if final temperatures were to exceed 6°C (11°F). It’s the roughly 10 percent chance of near-certain disaster that makes climate change costlier still.

Now we are truly in the realm of what Nassim Nicholas Taleb describes as a “Black Swan” and Donald Rumsfeld as “unknown unknowns.” We don’t know the full implications of an eventual 6°C (11°F) temperature change. We can’t know. It’s a blind planetary gamble. Devastating home fires, car crashes, and other personal catastrophes are almost always much less likely than 10 percent. And still, people take out insurance to cover against these remote possibilities, or are even required to do so by laws that hope to avoid pushing these costs onto society. Risks like this on a planetary scale should not—*must not*—be pushed onto society.

“Must not” is a strong phrase. It conjures images of bans or—in dollars and cents—infinite costs. That goes head-to-head against any economist’s idea of trade-offs. The costs of global warming may be high, perhaps higher than anyone thought possible. But surely, they can’t be infinite.

MONEY IS EVERYTHING

Trying to estimate the eventual temperature increase is one thing. But even if we knew that number with any precision for that one hot day in August 2100 in Phoenix, what actually concerns us isn’t necessarily how high temperatures climb. We care more about climate impacts, and how much they will cost society. Sea-level rise is one. Another is extreme events like droughts or hurricanes that might hit your home long before rising sea levels would drive you from it altogether.

The business of pinning down specific impacts is messy and fraught with its own uncertainties. There are known unknowns aplenty. Unknown unknowns may yet dominate. And tipping points and other nasty surprises seem to lurk around every corner. Some of them may put warming itself on overdrive. Releasing vast carbon deposits in Siberian or Canadian permafrost could prove to be a tipping point resulting in bad global warming feedbacks. Others may have relatively less influence on actual temperatures but have plenty of other impacts. Melting of Greenland and the West Antarctic ice sheets alone already raises sea levels by up to one centimeter (0.4 inches) each decade. If the Greenland ice sheet fully melted, sea levels would rise 7 meters (23 feet). Full melting of the West Antarctic ice

sheet would add another 3.3 meters (11 feet). That's not happening tomorrow or even this century. The IPCC's estimates of global average sea-level rise for this century top out at 1 meter (3 feet). But the tipping point at which the full eventual melting becomes inevitable will be passed much sooner. We may have already passed the tipping point for the West Antarctic ice sheet.

These compounding uncertainties—first from emissions to concentrations to temperatures, and then from temperatures to ultimate impacts measured in dollars and cents—make things extremely hard to get right. That hasn't stopped economists from trying.

One of the best is Bill Nordhaus. His DICE model—short for Dynamic Integrated Climate-Economy model—has been publicly available since the early 1990s. Generations of graduate students have played around with it, tried to poke holes in it, and derived estimates of “optimal” global climate policy. Nordhaus's own estimates of the social cost of carbon have been going up ever since the model was first released in 1992. Back then, his economically optimal response to climate change was a global carbon tax of about \$2 per ton of carbon dioxide (in 2014 dollars). That went hand in hand with global average warming climbing to 4°C (7.2°F) and beyond. In the tug-of-war between economic growth and a stable climate, growth won. Climate impacts have been catching up ever since, pushing unfettered, fossil-fueled growth further and further from being optimal. Today, Nordhaus's preferred “optimal” estimate is around \$20 per ton of carbon dioxide. The resulting final temperature increases now top out at around 3°C (5.4°F).

The search for the optimal carbon price is a hot button issue. Nordhaus's formally derived \$20 is lower still than

the average estimate of \$25 per ton, presented in his own book as an “illustrative” example. That, in turn, is lower than the current “central” U.S. government’s estimate of around \$40, derived from a combination of outputs from DICE and two other assessment models.

None of that yet factors in the proper cost of the tails, fat or otherwise. Nordhaus’s maximum average temperatures may stay below 3°C (5.4°F), but that’s the average. It still leaves unspecified the probability of topping 6°C (11°F) or more. Some other estimates attempt to take uncertainty more seriously. The U.S. government itself presents what it calls the “95th percentile estimate” as a proxy of sorts for capturing extreme outcomes. The optimal number there: over \$100 per ton of carbon dioxide emitted today.

What then does the central \$40 estimate include, and how is it derived? Two key issues loom large: dollar estimates of damages caused, and discounting. We’ll address them in turn.

HOW MUCH FOR A DEGREE OF WARMING?

Compare the average climates in Stockholm, Singapore, and San Francisco. Winters in Sweden are long, cold, and dark. You’ll have to wait for the summer months to get average highs above 20°C (68°F). Singaporeans don’t have this problem. Their average low is higher than Stockholm’s average high year-round. All that makes San Franciscans feel smug, fog and all. They enjoy stable Mediterranean climates year-round, with a week of rain in “winter.” Still, all three cities are thriving metropolises. Historians may even argue that all of them got their start because of winning geographies. What, then, should lead us to believe that one

climate is better or worse than another? Or that warmer average global temperatures come with costs?

The costs of climate change aren't the result of moving away from some mythical optimal climate. Stockholm may be a more pleasant place with a degree or two extra. Incidentally, that's precisely what Swedish scientist Svante Arrhenius, of greenhouse effect fame, suggested we may want to do deliberately: burn more coal "to enjoy ages with more equable and better climates, especially with regards to the colder regions of the earth." In Arrhenius's defense, he said so in 1908, after he had identified the greenhouse effect, but long before it became clear that there are significant costs to pumping carbon dioxide into the atmosphere. In the end, the costs of small temperature changes are, for the most part, the sum of the costs of changing what we've gotten used to. And it's not just that Swedes already own winter jackets and Singaporeans air conditioners. It's massive investments and industrial infrastructures, built around current climates—and current sea levels—that make temperature increases costly.

And once again, it isn't temperatures themselves that matter as much as what these rising temperatures entail. One such effect is rising sea levels. Then there are storm surges on top—by then stronger and more frequent precisely because of climate change. And all that's the perfectly "normal," *average* effect of sea-level rise baked into where we *are* already heading. None of this is yet taking into account fat tails or other catastrophic scenarios.

When models incorporate the latest science and quantify ever more of the damages likely to occur because of climate change, the estimated costs of carbon pollution go up. DICE & Co are perennially playing catch-up with the latest science. In 2010, the central U.S. government's

estimate of the social costs of one ton of carbon dioxide emitted in 2015 was around \$25. The 2013 iteration increased it to around \$40.



None of this is meant to decry the modeling efforts. Quite the opposite. Getting things right is incredibly difficult. If anything, it is a call to invest in economic modeling—in a big way. Nordhaus’s DICE model, as well as its main competitors, FUND and PAGE, were all started by one person, and have been painstakingly maintained, patched, and modified over years and decades by a small group of dedicated economists. Meanwhile, when big business tries to analyze what toothpaste flavor to sell where, it uses massive quantities of geo-spatial, customer-level data, analyzed by dozens of dedicated statisticians and programmers.

We certainly shouldn’t scrap economic climate models for their inadequacy. If anything we should be supercharging them: IBM-ifying their operation. There’s much more at stake here than with selling toothpaste. Yet Colgate and Procter & Gamble are competing with the help of massive data operations, while DICE can run on your home computer. More manpower and data would at least help the models incorporate the latest available information in real time.

Even if we did all of that, though, there would still be one major problem: How should we quantify the damages caused by potentially catastrophic climate change? More data won’t necessarily help us make inroads on that question.

DICE & Co mostly look to the past for guidance. Hundreds of scientific studies try to quantify the impacts of

global warming on anything from sea-level rise to crop yields to tropical storms to war. The task then is to translate those impacts into dollars and cents. We quickly run into two problems. For one, only a small part of known damages can be quantified. Lots are missing. The list of currently unquantified and—at least in part—unquantifiable damages spans everything from known respiratory illnesses from increased ozone pollution due to warmer surface temperatures to the effects of ocean acidification. Moreover, the only parts we can truly quantify are in a relatively narrow, low-temperature range: changes of fractions of a degree, maybe 1°C (1.8°F), or maybe 2°C (3.6°F) of global average warming. How can we estimate what happens at 5, 4, or even 3°C (9, 7.2, or 5.4°F)?

Extrapolate, extrapolate, extrapolate.

That's at least what current models do. Take what happens at 1 or 2°C and scale it up. We know that, because of tipping points and other possibly nasty surprises, we can't just look at things linearly. No one seriously proposes that. Instead, DICE mostly relies on something close to quadratic extrapolations: If 1°C causes \$10 worth of damages, then 2°C doesn't cause \$20—that's linear—but \$40. More specifically, Nordhaus estimates that warming of 1°C costs less than 0.5 percent of global GDP, 2°C costs around 1 percent, 4°C costs around 4 percent. Things take off after that, but even 6°C stays below 10 percent.

Mind you, that's a big absolute number: 10 percent of total, global economic output today would be around \$7 trillion. If they were to materialize, by the time these 6°C (11°F) changes were to hit a century or more from now, the fraction of damages would be multiplied by a large growth factor. But how can we be certain that it's the right number? We can't. Once we extrapolate damage estimates




Figure 3.3 Quadratic and exponential extrapolations of global economic damages

as far out as 6°C (11°F), it all becomes guesswork. Using a quadratic function is a convenient shortcut, but it's not much more than that. Lots of other extrapolations would fit the observed damages on the lower end of the scale but would yield wildly different results on the upper end. For instance, figure 3.3 shows how estimating exponential rather than quadratic warming yields starkly different results:

For 1°C and 4°C, the two lines are identical. For 2°C and 3°C, they are close enough to be indistinguishable given the uncertainties. At 5°C, things begin to diverge. By 6°C, they might as well be describing different planets. The quadratic extrapolation ends up at a bit under 10 percent of global economic output. The exponential comes in closer to 30 percent.

We aren't saying a 30 percent decline in output is any more correct than 10 percent decline should global average temperature increases hit 6°C (11°F). We just don't know. And no one else does either. One could tell stories about how 10 percent may be too high because people will be able to cope. Even with 6°C (11°F) of warming, Stockholm then will still be cooler than Singapore now. Or one could tell stories around how 30 percent may still be too low because neither Stockholm nor Singapore would be around to see the day. Their current coast lines would be on track to be submerged under several meters of water. *Would*, not *may*. But it's once again the deep-seated uncertainties around the eventual extent and timing of the consequences that add to the true costs.

|||||

It's also not at all clear that we should be thinking about damages as a percentage of output in any given year. Standard practice for DICE and other models is to assume that the economy hums along just fine until damages from climate change get subtracted at some point in the future. Catastrophic or not, conventional estimates of climate damages will feel small compared to the amazing increases in wealth that economic growth is assumed to bring. At a 3 percent annual growth rate, global economic output will increase almost twenty-fold in a hundred years. Subtracting 10 percent, 30 percent, or even 50 percent for climate damages after a hundred years will still leave the world many times richer than it is today. Climate change, in short, may be bad, but even the worst seems to leave the world much better off so long as economic growth remains robust.

Instead assume that damages affect output growth *rates* rather than output *levels*. Climate change clearly

affects labor productivity, especially in already hot (and poor) countries. Then the cumulative effects of damages could be much worse over time. That's the beauty—or here, the ugliness—of compound growth rates. All it took was a small but all-important change in a fundamental assumption.

Lastly, the way that climate damages are assumed to interact more generally with economic output matters a lot. DICE & Co assume that climate damages are a simple fraction of GDP: the higher the temperature, the greater that fraction. That seems like an innocuous enough assumption, but there are some stark implications. GDP and temperature just became interchangeable. Or rather: climate damages amounting to 1 percent of output can always be offset by a 1 percent increase in the output itself. *More GDP is good. If more GDP implies higher damages, increase GDP further and the world will still be better off.* It's in the DNA of many economists to make that assumption. Growth, after all, is generally good.

Alas, not all environmental damages can be offset so easily just by increasing GDP. Loss of human lives, ecosystems, or food aren't compensated so readily by increased consumer electronics. To put it in more stark terms, if the global food supply suffers from climate change, boosting GDP by building more iPhones won't do much for those who are starving. Coming up with better ways to produce food would. That's typically the rejoinder of those in favor of using the standard multiplicative model of damages. *Human ingenuity has seemingly outpaced environmental degradation in the past. Things always seem to be getting cheaper, smaller, faster, better. Technology will win the day once again.* Maybe.

But what if there are limits? What if we can't, at some point, substitute away from bad environmental outcomes

in one area by increasing output further? Then more GDP will no longer compensate so easily for worse climate damages. The usual logic around economic growth being able to make up for climate damages just got turned on its head: Richer societies tend to prefer a better environment more so than poorer ones. In this world, the *higher* we can expect future GDP to be, the more valuable it is to have done something about global warming pollution today.

According to one study, if we assume that damages are additive rather than multiplicative—that food and iPhones aren't interchangeable—the “optimal” global average temperature increase is cut in half. If the standard, multiplicative version leads to around 4°C (7°F) of optimal eventual warming, making the simple change to additive damages will result in a final optimal temperature increase of below 2°C (3.6°F). That's an enormous difference, and it goes to show the importance of the assumptions that are feeding into models like DICE & Co. “Garbage in, garbage out,” as the saying goes. Here it takes the form: “optimism in, optimism out.” Feed a slightly different functional form through the most standard of climate-economy models, and the optimal climate policy can look very different.

|||||

Once again, the inherent uncertainty is the biggest story. That goes for functional forms of the damage function as well as for lots of other factors. Even if we knew with certainty how emissions develop, how concentrations follow, how temperatures react, and how sea levels rise—and we don't—we would still need to translate it all into dollars and cents.

It's not useful to pick different types of extrapolations that deterministically project either 10 or 30 percent or more of economic damages by the time temperatures hit 6°C (11°F) above preindustrial levels. Rather, the correct approach is to do here what we just did for final temperature outcomes: look at the entire distribution of possible damages for each temperature outcome, not the expected damages conditional on any one temperature level. In other words, if temperatures were to go up 6°C (11°F), what's the probability of damages hitting 10 percent of GDP, or 30 percent, or any quantity in between or beyond? The problem is that we have no idea.

There's always a small chance that any particular final temperature wouldn't cause any damage. There's also always a small chance that it would cost the world. The most likely outcome may well be somewhere in the middle—maybe indeed somewhere in the 10 to 30 percent range for warming of 6°C (11°F)—but that's not the point. Or at least it's not sufficient. It's a "guesstimate" at best, a guess at worst.

Therefore, we simply can't give you another table the way we did for median temperature outcomes and probabilities of hitting 6°C (11°F). We don't know enough to fill in even the one row indicating average global damages at each temperature outcome. Bill Nordhaus's estimates around average expected damages—that warming of 1°C costs less than 0.5 percent of global GDP, 2°C costs around 1 percent, 4°C costs around 4 percent—could be a start. But even there, anything above around 2°C (3.6°F) is already largely guesswork. And we know much too little about the actual distribution of damages at each temperature level to estimate the third row for 50 percent or any other number for catastrophic impacts such as in table 3.2.

Table 3.2. Knowledge of economic damages decreases quickly with increased global average warming

<i>Final temperature change</i>	2°C (3.6°F)	2.5°C (4.5°F)	3°C (5.4°F)	3.5°C (6.3°F)	4°C (7.2°F)	4.5°C (8.1°F)	5°C (9°F)	5.5°C (10°F)	6°C (11°F)
Average global damages	1%	1.5%	2%	3%	4%	?	?	?	?*
Chance of damages >50% of economic output	?%	?%	?%	?%	?%	?%	?%	?%	?%

* Our range for average global damages from 6°C (11°F) of warming is 10 to 30 percent throughout the text, though that’s hardly scientific enough to merit mention in this table. It’s simply an extrapolation, using quadratic and exponential curves, from what we know—or think we know—happens at 1 or 2°C (1.8 or 3.6°F).

When it comes to high-temperature damages, the state-of-the-art economic models simply aren’t much better than fitting a curve around what we know at low temperatures, and extending it into what we don’t—well beyond the range of historically observed temperature increases into ones that mark uncharted territory for human civilization. Yet again, that’s not to decry these modeling efforts. It’s just to reiterate that inherent uncertainties will probably determine the final outcome.



All that begs the philosophical question: is some number better than no number at all? If Nordhaus’s estimate for the average global damages caused by final warming of 6°C (11°F) is 10 percent, and a simple exponential extrapolation gives us 30 percent, should we be using this 10–30 percent range at all? And what happens if we have fundamentally

mis-specified damages because, for example, they affect growth rates rather than output levels, or because damages are inherently additive rather than multiplicative?

But what's the alternative? If we didn't use these numbers in government benefit-cost analyses, we would essentially be accepting a climate damage estimate of zero. That's most definitely the wrong number. So better to go with the standard output of DICE and models like it. The U.S. government's \$40 per ton figure is as good as any in that regard, even if still a likely underestimate. Let's at least run with it for now to illustrate another important point.

HOW MUCH FOR A DEGREE OF WARMING ONE HUNDRED YEARS FROM NOW?

Whether damages at 6°C (11°F) of eventual warming are 10 percent or 30 percent of global economic output or not anywhere near that range may be anyone's guess. The one thing we know for sure is that we ought to discount whichever number we get. The basic logic of discounting is sound, and ever present: It's a combination of delayed gratification and risk. Having \$1 today is worth more than having it ten years from now. Answering the question of how much more seems to be as much an art as a science. But it doesn't need to be.

In fact, there's a website for it. Go to treasury.gov to find the interest rate for what are commonly viewed as the least risky investments imaginable: U.S. government bonds. Lend the United States of America \$100 today for up to thirty years, and see your investment grow each year at the rate displayed there. More specifically, you'd want to look

for Treasury Inflation-Protected Securities or TIPS. That way what you see is what you get in purchasing power. Inflation won't eat into your earnings. That rate has been hovering at around 2 percent a year for the past ten years. At the moment it's closer to 1 percent.

Contrast that range with the central estimate of 3 percent that the U.S. government uses in its social cost of carbon calculations. Nordhaus, in his DICE model, arrives at a default value of around 4 percent. Lord Nicholas Stern, in his *Stern Review on the Economics of Climate Change*, used 1.4 percent. He was also roundly criticized for that low choice at the time. So what is the right discount rate?

The short answer is that we don't know, but we are pretty sure the correct long-run discount rate should decline over time "toward its lowest possible value." That seems rather self-serving for anyone trying to argue for strong climate action today. A low discount rate implies that future climate damages will be more significant in today's dollars and, thus, favors strong climate action today. But there is, in fact, quite a bit of underlying science pointing in that direction. Once again, it's what we don't know that points mostly in one direction. The primary driver for low discount rates is uncertainty around the correct rate itself. Who knows what the discount rate should be a century or two from now? The less we know about the correct discount rate, the lower it ought to be. Since we know less about the right discount rate the further out we go, the rate ought to decline over time. So what exactly is that number? It's probably not 4 percent or 3 percent as currently used, but likely significantly lower, perhaps 2 percent or even less. It's so far in the future that we can't know for sure, but precautionary prudence dictates we should at least consider using low rates for long-term discounting.

Any of these numbers, though, focuses on risk-free rates. It's what you get for sure, no matter what happens to the world around you. The entire point of worrying about climate change was the unpredictability of it all. Should each potential future scenario then be discounted at the same rate?

CLIMATE FINANCE

We could do a lot worse than look to finance for cues on how to discount the uncertain future. When in doubt, ask those who actually stand to lose money from their decisions. Bob Litterman has spent most of his career at Goldman Sachs, serving in the late 1990s as head of firm-wide risk management before moving to asset management. He has lived and breathed the Capital Asset Pricing Model (CAPM) his entire life. In fact, he developed a variant, the Black-Litterman Global Asset Allocation Model. It allows for asset pricing decisions without making assumptions about expected returns for each type of asset. The less we know, the better his model performs against the standard version.

Litterman doesn't mince words when talking about the way some climate economists look at discounting: "They argue for high discount rates because of high opportunity costs for money, some estimate of the market return on capital. Waa!?! If that was the sole criterion, why would anyone invest in bonds, ever? We learned in finance why this is wrong sometime in the 1960s." In fact, CAPM was developed in the 1960s, and it has one simple premise: If an investment's fortunes rise in tough economic times, it will be more valuable than an identical investment that rises

and falls with the market. That link between its returns and the returns of the market is called “*beta*.” Low *beta* implies a weak link. A weak link increases the value of the investment.

That, in a sense, is the only reason why anyone would invest in government bonds that pay 1 or 2 percent rather than earn an expected 7 percent in the stock market. A high overall return is good, but it’s much less valuable if it pays off only in already good economic times—a high *beta*. U.S. government bonds have a low expected return but they also have a low *beta*. Many balanced investment portfolios include at least some bonds as a rainy-day fund for tough economic times.

This comes with stark implications for climate policy: If we somehow think that climate damages are small and will be worse when the economy is strong, discount rates should be higher. It will be fine to live through episodes of extreme weather events, say, because the storms won’t be all that bad, and they’ll hit only when GDP is high. This is one view supporting high discount rates. On the other hand, if we believe that climate damages will be large and go hand in hand with times when the economy is doing poorly, discount rates should be low. That might be a world in which climate change implies more extreme heat days, which in turn decrease labor productivity and, thus, GDP.

Or more directly, if there is around a 10 percent chance of a climate catastrophe that crashes economies and alters life as we know it, unless we change course, Finance 101 tells us that the discount rate on those damages far into the future should be low—maybe even lower than the 1 to 2 percent rate applied to assessing risk-free bonds. How low? Nobody knows for sure, but this is where we need to take a quick detour into Finance 102.

WALL STREET PUZZLES

Despite all its sophistication, modern finance leaves us with a number of fundamental puzzles. The “equity premium puzzle” tops that list. Investing in U.S. stocks returns, on average, 5 percent more than investing in U.S. government short-term bonds. This simple fact has haunted economists for decades. Standard economic models simply can’t replicate these basic facts. People ought not to be so risk-averse as to warrant that large a premium for investing in risky stocks. Yet they are. What gives?

Daily stock prices are among the most well-known facts. Newspapers print them. Comprehensive databases are freely available online. So questioning the underlying data won’t get us far. It’s also tough to see how we could blame it on laziness, biases, or some other human quirks that may or may not contribute to the puzzle. There’s a lot of money at stake, and most of it is managed by professionals, who should know better. The most natural place to look for the culprit then is economic theory itself. We know that each model simplifies reality. Do the standard models simplify too much for their own good?

It turns out that introducing potentially catastrophic risks to the standard models explains and even reverses the equity premium puzzle: fat tails in action. Market outcomes aren’t defined by the average fluctuations on a typical day. They are much more defined by what happens during extreme events, the kinds of things that should never happen but have since given us at least a week’s worth of “black” days in the past a hundred and fifty years: from *Black Monday* in October 1987 to *Black Friday* in September 1869 to an entire *Black Week* in October 2008. Taking these sorts of catastrophic risks more seriously justifies

large equity premiums, the amounts of money investors need to be paid to take the risk.

The same holds for climate risks. Potentially catastrophic climate events demand a “risk premium.” The higher the chance of these catastrophes, the more we ought to seek out the climate-equivalent of risk-free government bonds: avoiding carbon emissions in the first place.

There’s one more complicating factor to this story, and it comes back to discount rates and the all-important *beta*. The reason anyone invests in government bonds is because of their low *beta*, which makes them pay off in all states of the world, including the bad. Standard asset pricing models value these investments by assigning a low, sometimes even *negative* discount rate. The latter comes into play, for example, with contrarian short sellers that earn more when the stock market is lower.

That same insurance thinking ought to apply to climate damages, or rather to avoiding them. Bob Litterman, drawing the link to climate: “If the risk premium is large enough, then the insurance benefits could even require a negative discount rate and such a high current price of emissions that the price would actually be expected to drop over time as the problem diminishes and uncertainty is resolved.” The point seems blindingly obvious from an asset pricing perspective. It will come as a surprise to most focused on the proper way of discounting in the climate arena, where tying discount rates to “opportunity costs” or expected “market returns” is common practice—and where Lord Nicholas Stern’s 1.4 percent has long been considered the lower bound of what is an acceptable rate. But there is nothing magical about 1.4 percent, or about 1 percent. Even 0 percent doesn’t need to be a lower bound in theory. It isn’t to those shorting Wall Street. If investing in

a project pays more in the hardest of economic times, the proper discount rate may need to be lower than the lowest risk-free rate. We are not at all sure this is the relevant case for climate change, but it is a definite possibility, representing yet another big uncertainty.

With a large chance of catastrophe—a 10 percent chance of eventually hitting 6°C (11°F), say—and with that catastrophe associated with large economic costs—10 or 30 percent (or even much more) of global economic output—the proper treatment of these climate damages is to discount them at a rate maybe even lower than the risk-free rate of government bonds. As always, it's tough to pick a single number, but it makes it harder to argue for discount rates much above 1 or 2 percent.

TIMING IS EVERYTHING

We keep saying “eventual” in connection with warming of 6°C (11°F) and other such extreme scenarios, because any of these catastrophic temperature increases would play out over many decades and centuries. Large global average temperature increases won't happen tomorrow; nor would catastrophe strike overnight, at least not based on this calculation. In fact, the higher the final temperature increase, and the higher the chance of ultimate catastrophe, the longer it will take for both to materialize. That points to one of the more profound characteristics of climate change: its long-term nature. But it clearly doesn't mean that we can relax for the time being.

If a civilization-as-we-know-it-altering asteroid were hurtling toward Earth, scheduled to hit a decade hence, and it had, say, a 5 percent chance of striking the planet, we would surely pull out all the stops to try to deflect its path.

If we knew that same asteroid were hurtling toward Earth a century hence, we may spend a few more years arguing about the precise course of action, but here's what we wouldn't do: We wouldn't say that we should be able to solve the problem in at most a decade, so we can just sit back and relax for another 90 years. Nor would we try to bank on the fact that technologies will be that much better in 90 years, so we can probably do nothing for 91 or 92 years and we'd still be fine.

We'd act, and soon. Never mind that technologies will be getting better in the next 90 years, and never mind that we may find out more about the asteroid's precise path over the next 90 years that may be able to tell us that the chance of it hitting Earth is "only" 4 percent rather than the 5 percent we had assumed all along. That last point—increased certainty around the final impacts—is precisely where climate change has proven so vexing. Our estimate of the range of climate sensitivity isn't any more precise today than it was over three decades ago. And the chance of eventual climate catastrophe isn't 5 percent; our rough calculation based on IEA projections shows that it's likely closer to 10 percent or even more.

WHAT'S YOUR NUMBER?

Climate change is beset with deep-seated uncertainties on top of deep-seated uncertainties on top of still more deep-seated uncertainties. And that's just for going from emissions to concentrations to final temperatures. Further uncertainties prevent us from simply translating temperatures into economic damages, and none of that yet clarifies the uncertainties around the correct discount rates to calculate the optimal carbon price today. In each of these

steps, though, one thing is clear: because the extreme downside is so threatening, the burden of proof ought to be on those who argue that fat tails don't matter, that possible damages are low, and that discount rates ought to be high.

As little as we know about many of these uncertainties, we do know that the chance of eventual catastrophic warming of 6°C (11°F) or more isn't zero. It's slightly greater than around 10 percent under our conservative calibration.

Associated damages are anyone's guess, but we can only consider the implied "guess" of 10 percent of global economic output ventured by Bill Nordhaus's DICE model a lower bound. Following the same, admittedly imperfect logic could yield estimates anywhere from 10 percent to 30 percent or even well beyond. We don't know where in that range the true number is. We are pretty sure it's not less than 10 percent, and we do know that no one else knows the true number either. The most relevant question isn't whether expected damages at 6°C (11°F) are 10 or 30 percent of global economic output. The question ought to be: what is the full distribution of damages and what is the chance of significant economic collapse?

That leaves discounting, where at least we know that looking for an expected market return on capital to arrive at a discount rate of, say, 4 percent may be turning a blind eye on decades of asset pricing theory and practice. If we omit the rosy scenario where climate damages are small and will be worse when the economy is strong, we are looking at much lower rates than are currently bandied about. We don't know whether the right rate should be 2 percent, 1 percent, or even below. There may not be a single *climate beta*—the link between climate damages

and the general health of the economy—to justify using any one particular rate. But we can be pretty sure that the presence of big uncertainties around high final temperatures and catastrophic damages should drive discount rates down, not up. A rate of 2 percent might be our estimate for damages fifty years hence, and whatever rate it is, it ought to decline over time.

Where does all of that leave us? First, with the realization that it's easy to criticize. It's tougher to come up with a constructive alternative. Table 3.2, showing actual climate damages, is mostly blank for a reason, and it's not for lack of trying.

If the question is what single number to use as the optimal price of each ton of carbon dioxide pollution today, the answer should be: at least \$40 per ton of carbon dioxide, the U.S. government's current value. We know it's imperfect. We are pretty sure it's an underestimate; we are confident it's not an overestimate. It's also all we have. (And it's a lot higher than the prevailing price in most places that do have a carbon price right now—from California to the European Union. The sole exception is Sweden, where the price is upward of \$150. And even there, key industrial sectors are exempt.)

If the next question is how to decide on the proper climate policy, the answer is more complex than our rough benefit-cost analysis suggests. Pricing carbon at \$40 a ton is a start, but it's only that. Any benefit-cost analysis relies on a number of assumptions—perhaps too many—to truly come up with one single dollar estimate based on one representative model of something as large and uncertain as climate change.

Since we know that fat tails can dominate the final outcome, the decision criterion ought to focus on avoiding the

possibility of these kinds of catastrophic damages in the first place. Some call it a “precautionary principle”—better safe than sorry. Others call it a variant of “Pascal’s Wager”—why risk it, if the punishment is eternal damnation? We call it a “Dismal Dilemma”: while fat tails can dominate the analysis, how can we know the relevant probabilities of rare extreme scenarios that we have not previously observed and whose dynamics we understand only crudely at best? The true numbers are largely unknown and may simply be unknowable.



In the end, it’s risk management—existential risk management. And it comes with an ethical component. Precaution is a prudent stance when uncertainties about catastrophic risks are as dominant as they are here. Benefit-cost analysis is important, but it alone may be inadequate, simply because of the fuzziness involved with analyzing high-temperature impacts.

Climate change belongs to a rare category of situations where it’s extraordinarily difficult to put meaningful bounds on the extent of possible planetary damages. Focusing on getting precise estimates of the damages associated with eventual global average warming of 4°C (7°F), 5°C (9°F), or 6°C (11°F) misses the point. The appropriate price on carbon is one that will make us comfortable enough to know that we will *never* get to anything close to 6°C (11°F) and certain eventual catastrophe. *Never*, of course, is a strong word, since we know the chance of any of these temperatures happening even based on today’s atmospheric concentrations can’t be brought to zero.

One thing we know for sure is that a greater than 10 percent chance of eventual warming of 6°C (11°F) or

more—the end of the human adventure on this planet as we now know it—is too high. And that’s the path the planet is on at the moment. With the immense longevity of atmospheric carbon dioxide, “wait and see” would amount to nothing other than willful blindness.

Notes



PREFACE: POP QUIZ

Page ix—**Two quick questions:** Princeton’s Robert Socolow has started many a presentation with a version of this quiz, asking audiences whether they consider climate change “an urgent matter” and fossil fuels “hard to displace.” He groups the resulting views into four broad buckets, reproduced here with permission, and with slight modifications:

		<i>Is getting the world off fossil fuels difficult?</i>	
		No	Yes
<i>Is climate change an urgent problem?</i>	No	A low-carbon world unmotivated by climate considerations.	Perhaps most of the general public, and parts of the energy industry.
	Yes	Many environmentalists, including nuclear advocates.	Our working assumption.

Socolow, “Truths,” searches for solutions firmly grounded in this “working assumption.” Oliver Morton, editor at the *Economist*, introduced an August 2013 debate on geoengineering at the Massachusetts Institute of Technology with these two questions. Morton echoed Socolow’s conclusion that, to avoid cognitive dissonance, most people answer “Yes” to either one or the other question, but not both. In the packed lecture hall that evening at MIT, most answered “Yes” to both, a clear indication of the type of people currently attracted to geoengineering conversations.

Page x—**Standard economic treatments:** For a popular, standard perspective on the science and economics of climate change, see Nordhaus, *Climate Casino*. For more, see the “DICE” entry on page 36 in chapter 2.

Page 44—**Books:** Sunstein’s “Of Montreal and Kyoto,” and the later adaptation in Sunstein, *Worst-Case Scenarios*, provide a comparative history and analysis of the Montreal and Kyoto Protocols. He suggests a few reasons for why the former worked so well while the latter has at best led to small steps in the right direction. In particular, Sunstein makes a strong case that success of the one, and the failure of the other, had a lot to do with domestic benefit-cost analysis in the United States. For a terrific insider’s view on the making of the Montreal Protocol, see Benedick, *Ozone Diplomacy*. Barrett, *Environment and Statecraft*, uses, in part, the success of the Montreal Protocol to develop a theory on international environmental treaties, and what makes them work, or, in most cases, fail.

CHAPTER 3. FAT TAILS

Page 48—**more likely than not:** The IPCC attempts to assign plain-English terms to its consensus assessments: “more likely than not” corresponds to a likelihood of greater than 50 percent; “likely” corresponds to greater than 66 percent (not two-thirds; i.e. 67 percent); “very likely” corresponds to greater than 90 percent; “extremely likely” corresponds to greater than 95 percent. These terms were used to describe the likelihood of man-made global warming in, respectively, the IPCC’s Second, Third, Fourth, and Fifth Assessment Reports. According to Engber, “You’re Getting Warmer,” an early draft of the Fourth Assessment Report called for the highest category, “virtually certain,” which corresponds to greater than 99 percent probability, before settling at “very likely.” Engber discusses the history and implications of the IPCC’s probability assessments. For the latest formal guidance document, see Mastrandrea et al., “IPCC AR5 Guidance Note.” For more on the history and the scientific underpinnings, see Giles, “Scientific Uncertainty.” For a survey of how these probabilistic statements are perceived (and often misconstrued), see Budescu et al., “Interpretation of IPCC.”

Page 48—**decade ‘without warming’:** For more on the warming “pause” or “hiatus” of recent years, see “warmest in human history” on page 9 in chapter 1.

Page 48—**back to the 1800s:** See “Climate Science” on page 35 in chapter 2. For more on the history, and the future, see Roston, *The Carbon Age*.

Page 49—**Wally Broecker:** Broecker, “Climatic Change.”

Page 49—**climate sensitivity:** See “Climate Sensitivity” on page 35 in chapter 2.

Page 49—**well-established facts:** Stocker, “Closing Door,” and Matthews et al., “Proportionality of Global Warming,” are among the latest to discuss the proportional relationship between total warming and cumulative emissions.

Page 50—**400 ppm:** This is the concentration of carbon dioxide. Counting other greenhouse gases (without aerosols), concentrations are between 440 and 480 ppm of carbon dioxide–equivalent greenhouse gases, depending on the source. See “400 parts per million” on page 10 and also “2 ppm” on page 22 in chapter 1.

Page 50—**700 ppm:** See “700 ppm” on page 14 in chapter 1.

Page 50—**Ad Hoc Study Group:** Charney et al., “Carbon Dioxide and Climate.”

Page 50—**academic genius:** Gavin Schmidt tells the story on Real Climate.org, an excellent repository of the latest on climate change science (Schmidt and Rahmstorf, “11°C Warming”).

Page 50—**the “likely” range:** By 1990 the IPCC range was still 1.5–4.5°C (2.7–8.1°F). Ditto by 1995 and 2001. By 2007, the range narrowed somewhat, though in the wrong direction. It seemed that 1.5°C (2.7°F) was no longer in the cards. The new “likely” range was 2–4.5°C (3.6–8.1°F). By 2013, the most recent IPCC Assessment report, the range widened again right back to where it’s been all along: 1.5–4.5°C (2.7–8.1°F). For the relevant sections of the reports, see Working Group 1, chapter 5, of the *IPCC First Assessment Report*, Section B: Climate Modelling, Climate Prediction and Model Validation, of the *IPCC Climate Change 1992 Supplementary Report*, Working Group I of the *IPCC Second Assessment Report*, Working Group I of the *IPCC Third Assessment Report*, Working Group I of the *IPCC Fourth Assessment Report*, and Working Group I of the *IPCC Fifth Assessment Report*.

It is true that the confidence in the range has increased markedly over time. Specifically, “confidence today is much higher as a result of high quality and longer observational records with a clearer anthropogenic signal, better process understanding, more and better understood evidence from paleoclimate reconstructions, and better climate models with higher resolution that capture many more processes more realistically” (Working Group I of the *IPCC Fifth Assessment Report*, TFE.6; also see Box 12.2). Still, the IPCC chose to call the range “likely” (>66 percent confidence) rather than opt for a more certain assessment such as “very likely” (>90 percent).

Things may even be worse than before for another reason. In 1990, the IPCC ventured a “best guess” of 2.5°C (4.5°F) within the wider range. By 2007, the “most likely” quantity was 3°C (5.4°F). Not certainty, not even an actual “mean” or “median” in statistical terms,

but at least a single number—albeit a high one—around which to rally. By 2013, the IPCC issued no verdict as to which quantity would be most likely. That’s a step back in sureness. The IPCC did add other caveats, notably a less than 5 percent probability of climate sensitivity being below 1°C and a less than 10 percent probability of above 6°C. See “clearly more room” on page 51 as well as “more likely than not” on page 48 for a definition of the “likely” range itself.

Page 51—**defines “likely”**: See “more likely than not” on page 48 as well as prior note.

Page 51—**clearly more room**: The IPCC’s latest assessment report goes into a bit more detail: it describes anything below 1°C as “extremely unlikely” (0–5 percent) and anything above 6°C as “very unlikely” (0–10 percent) (Summary for Policymakers of Working Group I in the *IPCC Fifth Assessment Report*). The second row of this table translates the IPCC’s statements into actual probabilities for different climate sensitivities:

<i>Climate sensitivity</i>	<0°C	<1°C	<1.5°C	<>2.6°C	>3°C	>4.5°C	>6°C
IPCC (2013)	No data	0–5%	(“likely” between 1.5–4.5°C)			0–10%	
Our calibration	0%	1.7%	11%	50%	37%	11%	3.1%
			(78% probability of between 1.5–4.5°C)				

We calibrated a log-normal distribution by calculating an 11 percent probability of being greater than 4.5°C and an 11 percent probability of being below 1.5°C. Doing so interprets the IPCC’s numbers as conservatively as possible. The IPCC, for example, states that any figure above 6°C would be “very unlikely.” That implies a 0–10 percent range—5 percent, if we take a point estimate. However, if the IPCC authors wanted to say that it was, in fact, only 5 percent, they could have chosen to say “extremely unlikely.” By saying “very unlikely,” they, in effect, may have intended to ascribe a probability of between 5 and 10 percent—7.5 percent as the point estimate. Either way, our calibration arrives at a probability estimate of slightly over 3 percent for the chance of climate sensitivity being greater than 6°C, a “conservative” estimate for the purposes of our exercise that remains much below 7.5 percent.

Our interpretation of the “likely” range uses a similar logic: The IPCC definition of “likely” is between 66 and 100 percent. However, if the authors wanted to convey that the probability of being in the 1.5–4.5°C range was higher than 90 percent, they could have

chosen to call the range “very likely.” (In fact, “very likely” does have a firm definition in the guidance document for IPCC authors, while “extremely likely” is an additional term added by the authors involved in Working Group I of the *IPCC Fifth Assessment Report* (see Working Group I’s Summary for Policymakers). For comparison, see Mastrandrea et al., “IPCC AR5 Guidance Note.” Instead, the IPCC authors opted for the looser interpretation of “likely,” which leads us to believe that the true likelihood may not be between 66 and 100 percent but between 66 and 90 percent. We split the difference and use 78 percent with 11 percent probability of being below the likely range and 11 percent probability of being above the likely range.

Our median estimate is 2.6°C (3.9°F). The most commonly cited mean climate sensitivity, 3°C (5.8°F), is therefore closer to the two-thirds mark in our calibration: we assume a 63 percent probability of climate sensitivity being below 3°C (5.8°F) and, conversely, a 37 percent probability of climate sensitivity being above 3°C (5.8°F). The latter is the probability mentioned in the table. The remaining estimates are in the bottom row of the table above.

Recent papers in *Science* and *Nature* have made the argument that climate sensitivity is much more likely above 3°C than below it. Fasullo and Trenberth, “A Less Cloudy Future,” finds that models with lower climate sensitivities do not fully take into account albedo from changing cloud cover. Sherwood, Bony, and Dufresne, “Spread in Model,” takes this idea further, and posits that accurate accounting of the cloud mixing processes suggests a climate sensitivity greater than 3°C (5.8°F).

Page 51—**bottle of Champagne:** A lot of analysis has gone into discovering the delicate science behind a good bottle of champagne. The average 750 ml bottle contains 9 grams of carbon dioxide, and emits about five liters once opened (Liger-Belair, Polidori, and Jean-det, “Science of Champagne Bubbles”). Not to mention the 200,000 metric tons released each year transporting the bubbly around the world (Alderman, “A Greener Champagne Bottle”). One way to minimize the loss of dissolved carbon dioxide in your champagne is to pour it like you would a beer—down the side of the glass instead of right to the bottom. It may not be quite as sophisticated, but scientists assure it will make for a better taste (Liger-Belair et al., “Losses of Dissolved CO₂”). Ironically, the celebratory champagne might not taste as nicely as it would if we had no cause for celebration at all. Wine quality in the Champagne region of France is, in fact, expected to increase based on projected climate change (Jones et al., “Global Wine Quality”).

Page 52—**fat tail**: The technical definition of a “fat tail” is a distribution that approaches zero polynomially or slower. Conversely, the technical definition of a “thin tail” is a distribution that approaches zero exponentially or faster. A log-normal distribution, which we use, is in between thin and fat tailed. Some definitions call it “heavy tailed”: no longer thin, but not yet fat either. Log-normal distributions approach zero faster than polynomially but slower than exponentially. All that means that our calibration is more conservative than the IPCC’s own numbers, as it points to a chance of slightly above 3 percent of climate sensitivity being above 6°C. That compares to the IPCC’s stated “very unlikely” range of anywhere between 0 and 10 percent. See “clearly more room” on page 51.

Page 53—**scientific papers**: The conceptual starting point for calculations leading up to this table is Weitzman, “Modeling and Interpreting the Economics.” A further elaboration is found in Weitzman, “Fat-Tailed Uncertainty.” A version resembling this table, with calculations for three different probability distributions but based on the *IPCC Fourth Assessment Report*, appeared in Weitzman, “GHG Targets as Insurance.” The version here is based on fitting a log-normal probability distribution to climate sensitivity, as described in “clearly more room” on page 51:

<i>CO₂e</i> concentration (ppm)	400	450	500	550	600	650	700	750	800
Temperature increase for mean climate sensitivity (=3°C)	1.5°C	2.1°C	2.5°C	2.9°C	3.3°C	3.6°C	4.0°C	4.3°C	4.5°C
Temperature increase for median climate sensitivity (=2.6°C)	1.3°C	1.8°C	2.2°C	2.5°C	2.7°C	3.2°C	3.4°C	3.7°C	3.9°C
Chance of >6°C, given 2–4.5°C “likely” range (with 70% probability)	0.03%	0.3%	1.3%	3.3%	6.3%	10.2%	14.4%	19.2%	23.9%

CO ₂ e concentration (ppm)	400	450	500	550	600	650	700	750	800
Chance of >6°C, given 1.5–4.5°C “likely” range (with 70% probability)	0.2%	1.1%	2.8%	5.2%	8.1%	11.3%	14.6%	18.0%	21.3%
Chance of >6°C, given 1.5–4.5°C “likely” range (with 78% probability)	0.04%	0.3%	1.2%	2.7%	4.9%	7.6%	10.6%	13.9%	17.3%

Row 1 has concentrations of ultimate carbon dioxide–equivalent (CO₂e) concentrations. Row 2 shows final temperature increases based on these concentrations and an assumed climate sensitivity of 3°C, the median figure found from a calibration to the 2007 *IPCC Fourth Assessment Report* “likely” range of 2–4.5°C. Row 3 shows the temperature increase for an assumed climate sensitivity of 2.6°C, the median figure found for our “best” log-normal calibration to the 2013 *IPCC Fifth Assessment Report* “likely” range of 1.5–4.5°C. The latter is also what we present in the main text. Note that this single “most likely” temperature increase is below the average or expected figure. That’s because the distribution fitted around the IPCC “likely” range for climate sensitivity is assumed to be an asymmetric log-normal distribution, which cuts off at zero but has a long upward tail. Despite the uncertainties, no one would seriously argue that climate sensitivity should have a negative realization. The next three rows present various assumptions around the IPCC “likely” range: Row 4 assumes the old climate sensitivity range of 2–4.5°C, the recent consensus before the release of the *IPCC Fifth Assessment Report* in 2013. Row 5 widens the “likely” range to extend to 1.5°C on the lower end. Both rows assume a 70 percent probability that climate sensitivity is within the “likely” range, rounding up the IPCC’s 66 percent number for its definition of “likely.” The final row then splits the difference between 66 percent (“likely”) and 90 percent (“very likely”) to put a probability of 78 percent of being between 1.5 and 4.5°C and a probability of 11 percent of climate sensitivity being above 4.5°C. This works out to a probability of 3 percent that climate sensitivity is

above 6°C, which is conservatively low by almost any standards. The final row is what we present in the main text, with numbers further rounded for simplicity.

Page 55—**heavy**: See “fat tail” on page 52.

Page 55—**the median**: Note that we are using the median climate sensitivity of 2.6°C here to calculate temperature increase. Using the more conventional mean climate sensitivity of 3°C would translate into a (mean) temperature increase of 4.0°C based on concentrations reaching 700 ppm (rather than the median figure of 3.4°C presented in the main text).

Page 55—**Black Swan**: Taleb, *Black Swan*.

Page 55—**unknown unknowns**: Donald Rumsfeld popularized the term in the context of the U.S. invasion of Iraq, drawing the analogy at more than one occasion. The first mention was at a Pentagon news conference on February 12, 2002: “Reports that say that something hasn’t happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don’t know we don’t know. And if one looks throughout the history of our country and other free countries, it is the latter category that tend to be the difficult ones” (Morris, “Certainty of Donald Rumsfeld”). Rumsfeld later echoed the sentiment at least once more, at a NATO press conference on June 6, 2002 (Rumsfeld, “Press Conference”).

Economists typically credit Chicago economist Frank Knight for coming up with the idea (Knight, “Risk, Uncertainty, and Profit”). He made the technical distinction between “risk” and “uncertainty.” (Note that Knightian “risk” is different from what the average person—including the average scientist—calls “risk.” Any layman’s “existential risk,” including the way we use it in the text, is much closer to Knightian “uncertainty” than Knightian “risk.”) Richard Zeckhauser has added a third category: “ignorance.” Risk deals with known distributions. Uncertainty is not knowing which distribution to pick. Ignorance is when it’s unclear there even is a distribution. See Zeckhauser, “Unknown and Unknowable,” and a subsequent reaction: Summers, “Comments.”

Page 56—**bad global warming feedbacks**: Walter et al., “Methane Bubbling,” attempts to measure methane emissions from thaw lakes in Siberia, estimating that methane emissions from northern wetlands is 10–60 percent higher than previously thought. They find that the largest portion of the methane released comes from the thawing

permafrost around lake edges. This process is thought to be a critical one in previous times of climatic change. “Climate Science: Vast Costs of Arctic Change” estimates the total cost to society of methane released from thawing Siberian permafrost to be on the scale of \$60 trillion.

Page 56—**Melting of Greenland:** The Greenland ice sheet has a sea level equivalence of 7.36 m (24 feet), and full melting of the Antarctic ice sheet would mean a 58.3 m (191 feet) sea-level rise (see chapter 4, “Observations: Cryosphere,” in Working Group I of the *IPCC Fifth Assessment Report*). Full melting of the West Antarctic ice sheet by itself would lead to about 3.3 meters (11 feet) of sea-level rise (Bamber et al., “Potential Sea-Level Rise”). See “centuries of sea-level rise” on page 9 in chapter 1 for more on the irreversibility of it melting.

The *IPCC Fifth Assessment Report* found that the Greenland ice sheet has contributed on average 0.59 mm per year to sea-level rise from 2002 to 2011, while the Antarctica contribution is likely 0.4 mm per year for the same period. Both of these contribution rates have more than quadrupled from the average for 1992 to 2001. The observed global mean sea-level rise for the 1993 to 2010 period was 3.2 mm per year. The IPCC’s estimate for total sea-level rise under the worst-case scenario is 0.53 to 0.97 m (1.7 to 3.2 feet). The only situation they believe could increase sea level by 2100 significantly above this likely range would be if marine-based sections of the Antarctic ice sheet collapsed (see chapter 13, “Sea Level Change” in Working Group I of the *IPCC Fifth Assessment Report*).

Page 57—**DICE model:** See the “DICE” entry on page 36 in chapter 2.

Page 57—**\$2 per ton:** Nordhaus derives a number of \$5 per ton of carbon for 1990 to 1999 (Nordhaus, “Optimal Transition Path”). We convert this figure to dollars per ton of carbon dioxide and into 2014 dollars using the GDP deflator to arrive at \$2 per ton of carbon dioxide.

Page 57—**Nordhaus’s preferred “optimal” estimate:** Nordhaus, “Estimates of the Social Cost of Carbon,” presents a price of \$18.6 per ton of carbon dioxide emitted in 2015 (in 2005 dollars). Converted into 2014 dollars, the figure is around \$20 per ton. The paper presents both what Nordhaus considers the “optimal” path and various other scenarios, including one to keep global average temperature increases below 2°C. Note that this \$20 estimate is significantly higher than his “optimal” path derived only four years prior. Then the optimal figure for 2015 was \$12 (Nordhaus, “Economic Aspects”). Note also that the \$20 is lower than both Nordhaus’s set of “illustrative carbon prices needed for a 2½°C temperature limit” (figure 33 in

Nordhaus, *Climate Casino*) and the “central” estimate presented in the first table of the “Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866” for a ton of carbon dioxide emitted in 2015. The former is \$25 for a ton emitted in 2015. The latter is close to \$40 per ton, using an average of three models and a discount rate of 3 percent. Nordhaus’s own preferred discount rate is around 4.2 percent. He shows how the difference in discount rates explains most of the difference between the \$40 and his own \$20 estimate.

Page 58—**around \$40**: See “possibly much more” on page 23 in chapter 1.

Page 59—**more equable and better climates**: The full quote: “By the increasing percentage of [carbon dioxide] in the atmosphere, we may hope to enjoy ages with more equable and better climates, especially with regards to the colder regions of the earth, ages when the earth will bring forth much more abundant crops than at present, for the benefit of rapidly propagating mankind” (Arrhenius, *Worlds in the Making*, 63).

The general spirit of the importance of and opportunities in adapting to warmer climates is best represented in Kahn’s *Climatopolis*. There are surely costs, Kahn argues, but coping mechanisms create their own opportunities, especially for highly efficient cities.

Page 59—**playing catch-up**: See “\$2 per ton” and “Nordhaus’s preferred “optimal” estimate” on page 57 for a discussion of the evolution of Nordhaus’s DICE estimates. See “possibly much more” on page 23 in chapter 1 for a discussion of the U.S. government’s figures.

Page 60—**started by one person**: Bill Nordhaus created DICE. Richard Tol developed FUND, which is now largely maintained by David Anthoff: <http://www.fund-model.org/>. Chris Hope was the driving force behind PAGE: <http://climatecolab.org/resources/-/wiki/Main/PAGE>.

Page 60—**massive data operations**: The underlying global circulation models used by climate scientists and feeding into the IPCC reports are indeed computationally complex. However, integrated assessment models then rely on much-simplified output, in DICE’s case on the Model for the Assessment of Greenhouse Gas Induced Climate Change (MAGICC), which in turn is a much-simplified version of underlying climate models. DICE itself is freely available on Bill Nordhaus’s website and even runs in Excel: <http://www.econ.yale.edu/~nordhaus/>.

Page 61—**Lots are missing**: See, for example, Howard, “Omitted Damages.” Van den Bergh and Botzen, “Lower Bound,” similarly present

climate change effects that are inadequately captured by models like DICE. Most of these effects would increase estimates of the social cost of carbon. Some may also decrease it. (See “possibly much more” on page 23 in chapter 1.)

Page 61—**quadratic extrapolations**: DICE uses an inverse quadratic loss function linked to temperature (T), where loss is defined as equal to $1/[1+aT+bT^2]$.

Page 62—**as far out as 6°C**: In fact, Nordhaus, *Climate Casino*, cuts off the graph at 5°C (9°F), implying that any damages due to temperature changes beyond that level are much too uncertain (or perhaps rare) to contemplate.

Page 63—**damages affect output growth rates**: See Pindyck, “Climate Change Policy,” and Heal and Park, “Feeling the Heat,” who link temperature to labor productivity via human physiology. They find high temperatures decrease productivity in already hot—and often poor—countries, while higher-than-average temperatures increase productivity by similar amounts in cool—and typically rich—countries.

Moyer et al., “Climate Impacts on Economic Growth,” similarly show large impacts on the Social Cost of Carbon estimate of changing the climate impact from output levels to productivity: “even a modest impact of this type increases SCC estimates by many orders of magnitude.”

Page 65—**damages are additive**: See Weitzman, “Damages Function.” For a complementary take—focusing on the idea of “relative prices”—see Sterner and Persson, “Even Sterner Review.” The fundamental distinction between multiplicative versus additive damage functions rests on questions of substitutability. The (implicit) assumption for multiplicative damages is unit substitutability between economic sectors and environmental amenities within the utility function. Additive damages assume less (to no) substitutability across these sectors in the utility function.

Page 68—**likely underestimate**: See “possibly much more” on page 23 in chapter 1.

Page 68—**discount whichever number**: Many a book and article has been written on the topic. Gollier, *Pricing the Planet's Future*, ranks among the best general introductions.

Page 68—**is worth more**: In fact, having one dollar today is typically worth a lot more than having it tomorrow. Whereas the same one-day difference a hundred years from now is barely noticeable. From today's perspective, a hundred years plus one day is pretty much the same as a hundred years. Quite naturally, humans tend to discount

the first day more heavily than that one day a hundred years from now. The technical term for this particular phenomenon is “hyperbolic discounting,” most prominently introduced to economics in Laibson, “Golden Eggs.”

Page 69—**2 percent a year:** “10-Year Treasury Inflation-Indexed Security, Constant Maturity.”

Page 69—**criticized for that low choice:** Weitzman, “A Review,” argues that “the *Stern Review* may well be right for the wrong reasons,” the low discount rate being one of the wrong reasons.

Page 69—**decline over time:** See Weitzman, “Gamma Discounting,” for the declining discount rate numbers mentioned in the text. For a later consensus view around the logic behind declining discount rates, not necessarily the specific numbers, see Arrow et al., “Determining Benefits and Costs,” and Cropper et al., “Declining Discount Rates.” France and the United Kingdom, for example, use declining discount rates, but they do not agree on the exact rate: France’s starts at 4 percent and declines to a bit over 2 percent for numbers 300 years in the future; the UK’s starts at 3.5 percent and declines to 1 percent after 300 years.

For an application that reconciles some important technical differences in the application of the basic logic, see Gollier and Weitzman, “Distant Future.” It concludes that “The long run discount rate declines over time toward its lowest possible value.”

To see the reason behind this declining rate, consider the following thought experiment: Assume we don’t know whether the true discount rate for damages a hundred years from now should be 1 percent or 7 percent. The former is on the lower end of rates for U.S. Treasury bills, which come as close to a risk-free investment as possible. The latter comes from the obscure but all-important “Circular A-94,” which the powerful U.S. government’s Office of Management and Budget suggests as the base-case analysis for all government investment and regulatory decisions (OMB, “Circular No. A-94 Revised”). Note that the 7 percent rate is not a risk-free rate in any sense of that word. In fact, it deliberately deals with risky investment decisions. The big question there is whether the rate for riskier investments ought to go up or down, something discussed in more detail in the text itself.

OMB’s other, and more appropriate, rate for something as far out as a hundred years is 3 percent. That’s also the government’s base case for the \$40 social cost of carbon. But for argument’s sake, let’s use the 7 percent figure for now. It’s certainly an upper bound

of sorts. Hardly anyone would sensibly argue for a higher discount rate. There's a good reason why: \$100 a hundred years from now, discounted at 7 percent is worth 9 cents today. Invest less than a dime today at a rate of return of 7 percent, and expect to get \$100 a century hence. For an investor, that's not bad. For discounting climate damages a hundred years out, it makes them almost worthless today. That, of course, is the exact line of reasoning some use to argue why climate damages don't matter all that much. Why worry about costs of global warming in a century, if all it takes is setting aside relatively little money today to cover the damages? All that holds true at 7 percent. But \$100 a hundred years from now, discounted at 1 percent, is worth \$37 today. That's quite a bit more.

Let's split the difference and use a 4 percent discount rate, halfway between 1 and 7 percent: now \$100 a hundred years from now is worth \$1.8 today. That's much, much closer to 9 cents from the 7 percent discount rate than the \$37 from a 1 percent rate. But that's only one way of "splitting the difference." What if we just didn't know whether the rate should be 1 percent or 7 percent?

Put the chance of either rate at 50–50. That's a 50 percent probability that the correct number should be 9 cents, and a 50 percent probability that it should be \$37. On average, that's roughly \$18. That average of the discounted numbers is much higher than the number using the average discount rate of 4 percent = $(7 \text{ percent} + 1 \text{ percent})/2$. In fact, in our example, the difference is a factor of ten: \$18 versus \$1.80. And the difference increases the further out you go.

Lastly, for a different argument for declining discount rates, see Heal and Millner, "Agreeing to Disagree." They suggest that the choice of discount rate is an "ethical primitive" to arrive at the same conclusion of declining discount rates.

Page 70—**Black-Litterman Global Asset Allocation Model:** Black and Litterman, "Global Portfolio Optimization."

Page 71—**stark implications:** "If climate risk dominates economic growth risk because there are enough potential scenarios with catastrophic damages, then the appropriate discount rate for emissions investments is lower than the risk-free rate and the current price of carbon dioxide emissions should be higher. In those scenarios, the 'beta' of climate risk is a large negative number and emissions mitigation investments provide insurance benefits. If, on the other hand, growth risk is always dominant because catastrophic damages are essentially impossible and minor climate damages are more likely to occur when growth is strong, times are good, and marginal utility

is low, then the ‘beta’ of climate risk is positive, the discount rate should be higher than the risk-free rate, and the price of carbon dioxide emissions should be lower” (Litterman, “Right Price”). For earlier use of “*beta*” in the climate context, see Sandsmark and Vennemo, “Portfolio Approach,” for an argument for a negative *beta* of mitigation investments. Gollier, “Evaluation of Long-Dated Investments,” makes the case for a positive *beta*.

Page 72—**equity premium puzzle**: For an overview, see Mehra, “Equity Premium Puzzle.”

Page 72—**reverses the equity premium puzzle**: For a technical exploration of this argument, see Weitzman, “Subjective Expectations,” or Barro, “Rare Disasters.” See Mehra, “Equity Premium Puzzle,” for alternative explanations of the equity premium puzzle and a survey of the ongoing debate.

Page 72—**“black” days**: *Black Monday* on October 19, 1987, saw the Dow drop 22 percent; *Black Tuesday* on October 29, 1929, marked the beginning of the Great Depression; *Black Wednesday* on September 16, 1992, earned George Soros a billion pounds betting against the Bank of England; *Black Thursday* on October 24, 1929, saw Wall Street lose over 10 percent almost at the opening bell (recall *Black Tuesday* just above for what happened next); *Black Friday* on September 24, 1869, saw markets crash after a failed attempt to corner the gold market. None of it should be confused with the *Black Week* beginning Monday, October 6, 2008, when the Dow fell 18 percent by Friday. The front-page articles of the *Wall Street Journal* after each of these events make for interesting, contemporaneous reading: For the reaction to what we now know as “Black Monday,” see Metz et al., “Crash of ’87.” After Black Tuesday and Thursday, the two days that are seen as the beginning of the Great Depression, the WSJ seems surprisingly nonchalant. “Pressure Continues: Stocks Sink Lower under Record Volume of Liquidation,” (published October 30, 1929) recognizes the huge drop in stock prices from the day before, but also states that “industrial activity is on a large scale and sound basis with no real indications of a depression in prospect,” and projects that “after the initial shock has worn off the decline will prove beneficial in many ways by releasing funds from market to industry.” “Demoralized Trading: Stocks Break on Record Volume—Banking Support Starts Rally” (published October 25, 1929), somewhat more awed by the situation, opens with the statement, “Yesterday’s market was in many respects the most extraordinary in the history of the Stock Exchange.” However, the article also ends with a projection that the market would turn around

soon. Zweig, “What History Tells Us,” looks at the preceding “Black Week” in comparison to and in the context of the “Great Crash.” For the front-page article of London’s *Financial Times* on the day after Black Wednesday, see Stephens, “Major Puts ERM Membership on Indefinite Hold.”

Page 73—**drawing the link:** Litterman, “Right Price.”

Page 74—**won’t happen tomorrow:** We have a much better idea of warming in the short and medium term: For the next two decades (2016 to 2035), the Summary for Policymakers in Working Group I of the *IPCC Fifth Assessment Report* finds, with “medium confidence,” a “likely” additional warming of between 0.3 and 0.7°C (0.5 and 1.3°F) relative to the past two decades (1986–2005).

For the final two decades of this century, the predictions diverge dramatically. Depending on which scenario one chooses, average global warming relative to the past two decades could be anywhere from 0.3–1.7°C, to 2.6–4.8°C (0.5–3.1°F, to 4.7–8.6°F), an enormous range with dramatically different consequences. And those are just the likely ranges. See, for example, “0.3 to 1 meters” on page 5 of chapter 1 for the implications for sea-level rise.

Note that all these estimates, including for sea-level rise, are relative to the two decades ending in 2005. 4.8°C (8.6°F) of additional warming relative to “today” would mean total warming of 5.5°C (9.9°F) from preindustrial levels.

Page 74—**the longer it will take:** See Roe and Bauman, “Climate Sensitivity,” for this point. They employ a standard willingness-to-pay framework to conclude that fat tails may not be that costly. (For a contrasting conclusion by the same [lead] author, see Roe, “Costing the Earth.”)

Page 77—**European Union:** See Ellerman, Convery, and de Perthuis’s *Pricing Carbon* for an early yet comprehensive survey of the EU’s emissions trading system.

Page 77—**sole exception is Sweden:** See Hammar, Sterner, and Åkerfeldt, “Sweden’s CO₂ Tax,” and Johansson, “Economic Instruments in Practice.”

Page 77—**decision criterion:** For a recent elaboration on the point around alternative decision criteria, see Heal and Millner, “Uncertainty and Decision.” Also see Millner, Dietz, and Heal, “Scientific Ambiguity and Climate Policy.”

Page 78—**ethical component:** For a climate scientist making the strong moral case, see Roe, “Costing the Earth.” For a moral philosopher making the strong case for economists to engage on the moral dimension, see Sandel, “Market Reasoning as Moral Reasoning.”