CHAPTER 1

911

Thank Russian police corruption for footage that eluded NASA and every other space agency. On February 15, 2013, an asteroid as wide as 20 meters (66 feet) exploded in the sky above the Russian city of Chelyabinsk during the morning commute hours, causing a blast brighter than the sun. It didn't take long for some spectacular videos to appear online, mostly from dashboard cameras many Russian drivers have to protect themselves against the whims of traffic cops. The blast injured 1,500, most because of glass shattered by the explosion. It was a sobering wakeup call for space agencies to ramp up their asteroid detection and defense capabilities.

The money for such efforts is perennially in short supply. But the technical means are there, or at least they could be. A U.S. National Academy study estimates it would take ten years and around $2 or 3 billion to launch a test to deflect an asteroid bound to hit Earth. It may not be as glamorous as sending a man to the moon within the decade, but it may be at least as important.

While the Chelyabinsk asteroid would have been too small to deflect, it would have still been nice to know about it in advance. The chance of a larger asteroid hitting us is small, but it's there. Educated guesses put it as a 1-in-1,000-year event. That's a 10 percent chance each century. We haven't yet spent the money to know for sure. The fact, though, is that a few billion dollars would allow

NASA and others both to catalogue the hazards and to defend against them. That's a small amount when measured against the costs of a potentially civilization-destroying threat. Around 65 million years ago it was a giant asteroid that caused the globe's fifth major extinction event, killing the dinosaurs.

Climate change isn't exactly hurtling toward us through outer space. It's entirely homegrown. But the potential devastation is just as real. Elizabeth Kolbert argues convincingly based on her book *The Sixth Extinction* how this time around: "We are the asteroid." In fact, by one recent scientific assessment, we are slated to experience global changes at rates that are at least ten times faster than at any point in the past 65 million years.

As Hurricane Sandy was whipping the Eastern Seaboard, leaving Manhattan below the Empire State Building partially flooded and almost entirely without power, New York governor Andrew Cuomo wryly told President Barack Obama that: "We have a 100-year flood every two years now." Hurricane Irene in August 2011 caused the first-ever preemptive weather-related shutdown of the entire, century-old New York City subway and bus system. It took only fourteen months for the second shutdown. Sandy hit in October 2012. All told, Irene killed 49 and displaced over 2.3 million. Sandy killed 147 and displaced 375,000.

New York, of course, is far from unique here. Typhoon Haiyan slammed the Philippines in November 2013, killing at least 6,000 people and displacing four million. Not even a year earlier, Typhoon Bopha struck the country, killing over a thousand and displacing 1.8 million. The

European summer heat wave in 2003 killed 15,000 in France alone, over 70,000 in Europe. The list goes on, spanning both poor and rich countries and continents.

Society as a whole—especially in rich places like the United States and Europe—has never been as well equipped to cope with these catastrophes as it is today. As is so often the case, the poor suffer the most. That makes these recent deaths and displacements in places like New York all the more remarkable.

What likens these storms and other extreme climatic events to asteroids is that they both can be costly, in dollars and in deaths. The important and clear differences show that the climate problem is costlier still.

First the obvious: Major storms have hit long before humans started adding carbon dioxide to the atmosphere. However, warmer average temperatures imply more energy in the atmosphere implies more extreme storms, floods, and droughts. The waters off the coast of New York were 3°C (5.4°F) warmer than average during the days before Sandy. The waters off the coast of the Philippines were 3°C (5.4°F) warmer than average just as Haiyan was intensifying on its path to make landfall. Coincidence? Perhaps. The increase off New York happened at the surface. The increase off the Philippines happened 100 meters (330 feet) below. But the burden of proof seems to rest on those questioning the link from higher temperatures to more intense storms.

That's particularly true, since the best research goes much beyond drawing circumstantial links. The science isn't settled yet, but the latest research suggests that climate change will lead both to more *and* bigger storms. Though hurricanes are among the toughest climatic events to link directly to climate change, mainly because of how rare they are. It's easier to draw the direct link from climate change to more common events like extreme temperatures, floods, and droughts.

Think of it like drunk driving: Drinking increases the chance of a car crash, but plenty of crashes happen without elevated blood alcohol levels. Or liken it to doping in sports: No single Barry Bonds home run or Lance Armstrong Tour de France stage win can be attributed to doping, nor did doping act alone. Bonds still had to hit the ball, and Armstrong still had to pedal. But doping surely helped them hit farther and bike faster. Major storms, like home run records and multiple *Le Tour* wins, have happened before. None of that means steroids or elevated levels of red blood cells in an athlete's blood had no effect. Something similar holds for elevated levels of carbon dioxide in the atmosphere.

Researchers are getting increasingly better at using "attribution science" to identify the human footprint even in single events. The UK's National Weather Service, more commonly known as the Met Office, has a Climate Monitoring and Attribution team churning out studies that do just that. One such study found with 90 percent confidence that "human influence has at least doubled the risk of a heatwave exceeding [a] threshold magnitude" of mean summer temperature that was met in Europe in 2003, and in no other year since 1851. Links will only become clearer in the future, both because the science is getting better and because extreme weather events are becoming ever more extreme.

Governor Cuomo's "100-year flood every two years" comment may have been a throw-away line, but he was on to something. By the end of the century, we can expect today's 100-year flood to hit as frequently as once every three to twenty years. That's a century out, long after our lifetimes, but we know that we can't wait that long to act. Already, the annual chance of storm waters breaching Manhattan seawalls has increased from around 1 percent in the 19th century to 20 to 25 percent today. That means lower Manhattan can expect some amount of flooding every four to five years.

Unlike with asteroids, there's no $2-to-3-billion, ten-year NASA program to avoid the impact of storms and other extreme climatic events like floods and droughts. Nor is there a quick fix for less dramatic events like the ever faster rising seas. As a first line of defense, higher seawalls would surely help. But they can go only so far for so long. Higher seas make storm surges all the more powerful, and higher seas themselves come with plenty of costs of their own. Imagine standing in the harbor of your favorite coastal city. Then imagine standing there at the end of the century with sea levels having risen by 0.3 to 1 meters (1 to 3 feet). It will only be a matter of time before higher seawalls won't do, when the only option will be retreat.

By then, it will be too late to act. We can't re-create glaciers and polar ice caps, at least not in human timescales. The severity of the problems will have been locked in by past action, or lack thereof. Future generations will be largely powerless against their own fate.

One possible response that attempts to provide a quick fix is large-scale geoengineering: shooting small reflective particles into the stratosphere in an attempt to cool the planet. Geoengineering is far from perfect. It comes with lots of potential side effects, and it's no replacement for decreasing emissions in the first place. Still, it may be a useful, temporary complement to more fundamental measures.

(We will start exploring the full implications of geoengineering in chapter 5.)

None of what we've talked about thus far even deals with the true worst-case scenarios. Having the climatic equivalent of ever more Chelyabinsk-like asteroids hit us is bad, but there are ways to cope. For relatively small asteroids, it's seeking shelter and moving away from windows. For relatively small climatic changes, it's moving to slightly cooler climates and higher shores. That's often easier said than done, but at least it's doable. For much more dramatic climatic consequences—such as a crippling of the world's productive agricultural lands—it's tough to imagine how we'd cope in a way that wouldn't cause serious disruptions.

Meanwhile, standard economic models don't include much of this thinking. Many observers regard average global warming of greater than 2°C (3.6°F) above preindustrial levels as having the potential to trigger events deserving of various shades of the label "catastrophe." Economists typically have a hard time making sense of that term. They need dollar figures. Does a catastrophe then cost 10 percent of global economic output? 50 percent? More?

While it's indeed necessary to translate impacts into dollars and cents, such benefit-cost analyses can act as only one guide for how society ought to respond. We should also take into account the potential for planet-as-we-know-it-altering changes in the first place. First and foremost, climate change is a risk management problem—a catastrophic risk management problem on a planetary scale, to be more precise.

CAMELS IN CANADA

If one wanted to imagine an all but intractable public policy problem, climate change would be pretty close to the ideal. Today's storms, floods, and wildfires notwithstanding, the worst effects of global warming will be felt long after our lifetimes, likely in the most unpredictable of ways. Climate change is unlike any other environmental problem, really unlike any other public policy problem. It's almost uniquely *global*, uniquely *long-term*, uniquely *irreversible*, and uniquely *uncertain*—certainly unique in the combination of all four.

These four factors, call them the *Big Four*, are what make climate change so difficult to solve. So difficult that—short of a major jolt of the global, collective conscience—it may well prove too difficult to tackle climate change just by decreasing emissions and adjusting to some of the already unavoidable consequences. At the very least we'll need to add *suffering* to the list. The rich will adapt. The poor will suffer.

Then there's the almost inevitable-sounding geoengineering, attempting a global-scale techno fix for a seemingly intractable problem. The most prominent geoengineering idea would have us deliver tiny sulfur-based particles into the stratosphere in an attempt to engineer an artificial sun shield of sorts to help cool the planet.

Everything we know about the economics of climate change seems to point us in that direction. Geoengineering is so cheap to do crudely, and it has such high leverage, that it almost has the exact opposite properties of carbon pollution. It's the "free-rider" effect of carbon pollution that has caused the problem: it's in no one's narrow self-interest to do enough. It's the "free-driver" effect that may push us to geoengineer our way out of it: it's so cheap that someone will surely do it based on their own self-interest, broader consequences be damned.

But let's not go there quite yet. Let's first tackle the Big Four in turn, beginning with why climate change is the ultimate "free-rider" problem:

Climate change is uniquely global. Beijing's smog is bad. So bad, that it comes with real and dramatic health effects that have prompted city officials to close schools and take other drastic actions. But Beijing's smog—or that in Mexico City or Los Angeles, for that matter—is mostly confined to the city. Chinese soot may register at measuring stations on the U.S. West Coast, much like Saharan dust may on occasion blow to central Europe. But all these effects are still regional.

That's not true for carbon dioxide. It doesn't matter where on the planet a ton is being emitted. Impacts may be regional, but the phenomenon is global and—among environmental problems—almost uniquely so. The ozone hole over the Antarctic is bad, but even at its height it has never reached the level of engulfing the globe. The same goes, say, for biodiversity loss or deforestation. These are regional problems. It's climate change that ties them together into phenomena with global implications.

The global nature of global warming is also strike one against enacting sensible climate policy. It's tough enough to get voters to enact pollution limits on themselves, when those limits benefit them and only them, and when the benefits of action outweigh the costs. It's a whole lot tougher to get voters to enact pollution limits on themselves if the costs are felt domestically but the benefits are global: a planetary "free-rider" problem.

Climate change is uniquely long-term. The past decade was the warmest in human history. The one before was the second-warmest. The one before that was the third-warmest. "Americans are noticing changes all around them," as the 2014 U.S. National Climate Assessment puts it. Changes are nowhere as evident as above the Arctic Circle: Arctic sea ice has lost half of its area and three-quarters of its volume in only the past thirty years. The *Foreign Policy* article describing "The Coming Arctic Boom" takes all of this as given. Then there are the visible changes all around. Again, from the National Climate Assessment: "Residents of some coastal cities see their streets flood more regularly during storms and high tides. Inland cities near large rivers also experience more flooding, especially in the Midwest and Northeast. Insurance rates are rising in some vulnerable locations, and insurance is no longer available in others. Hotter and drier weather and earlier snowmelt mean that wildfires in the West start earlier in the spring, last later into the fall, and burn more acreage." Climate change is here, and it's here to stay.

None of that should mask the fact that most of the worst consequences of climate change are still remote, often caged in global, long-term averages: global average surface temperature projections for 2100, or global average sea level projections for decades and centuries out. Strike two against sensible climate policy: the worst effects are far off—never mind that avoiding these predictions would entail acting now.

Climate change is uniquely irreversible. Even if we stopped emitting carbon tomorrow, we would have decades of warming and centuries of sea-level rise locked in. The eventual, full melting of large West Antarctic ice sheets may already be unstoppable. More extreme weather events are already here and will be with us for some time to come.

Over two-thirds of the excess carbon dioxide in the atmosphere that wasn't there when humans started burning coal will still be present a hundred years from now. Well over one-third will still be there in 1,000 years. These changes are long-term, and—at least in human timescales—virtually irreversible. Strike three.

As if three strikes weren't enough, there's another unique characteristic of climate change to round out the Big Four, and it may be the biggest one of them all: *uncertainty*—everything we know that we don't know, and perhaps more importantly, what we don't yet know we don't know.

Last time concentrations of carbon dioxide were as high as they are today, at 400 parts per million (ppm), the geological clock read "Pliocene." That was over three million years ago, when natural variations, not cars and factories, were responsible for the extra carbon in the air. Global average temperatures were around 1–2.5°C (1.8 to 4.5°F) warmer than today, sea levels were up to 20 meters (66 feet) higher, and camels lived in Canada.

We wouldn't expect any of these dramatic changes *today*. The greenhouse effect needs decades to centuries to come into full force. Despite the recent changes in the Arctic, ice sheets need decades to centuries to melt. Global sea levels take decades to centuries to adjust accordingly. Carbon dioxide concentrations may have been at 400 ppm three million years ago, whereas rising sea levels lagged decades or centuries behind. That time difference is important and points to the long-term nature and irreversibility of it all. See strikes two and three.

But all that's small consolation, and there's an important twist to strike four.

DEEP UNCERTAINTIES

The best available climate models come close in their temperature projections to what the world experienced during the Pliocene, but they aren't predicting sea levels of 20 meters (66 feet) higher. Nor do they predict camels wandering around Canada. Not now. Not hundreds of years from now. That's true for two important reasons.

First, most climate models are unduly skewed toward the known, sometimes making them much too conservative. Until recently, most climate models predicted rising sea levels only based on thermal expansion of the oceans (and the melting of mountain glaciers), but they did not include the effects of melting ice sheets. Warmer waters take up more space, leading to higher sea levels. That mechanism alone has indeed contributed to over a third of sea-level rise in the past two decades. It's also clear that melting glaciers in Greenland and Antarctica raise sea levels, but by how much is highly uncertain. Call it a "known unknown." Until recently, scientific understanding of melting polar ice caps had been so poor that most models simply left it out.

Second, even though climate models do get a lot of things right, there are fundamental things that we don't understand about the way the climate works. The averages are bad enough. While 0.1°C (0.2°F) of average global surface warming per decade sounds rather manageable and perhaps even pleasant, few dispute that a century or more of warming at this pace would lead to serious costs. But these averages hide two distinct sets of uncertainties that could pose the real problems.

The first set of uncertainties is inherent in any kind of global, long-term estimate. Presenting just the global average numbers masks at least four important facts: First, temperatures in the past century have been increasing at an increasing rate. Second, despite that generally increasing trend, temperatures fluctuate across years and decades. (Hence the infamous "decade without warming.") Third, air over the oceans is usually cooler than over land. Since two-thirds of the world is ocean, a global average increase of 0.07°C (0.13°F) per decade translated to about a 0.11°C (0.20°F) increase over land. Finally, temperatures over the poles have warmed more than elsewhere. Arctic temperatures are expected to increase at a rate more than twice the global average. That's particularly bad, since the poles are also where most of the world's remaining ice is. Melting ice on land above sea level means higher seas, as the latest sea-level projections now officially acknowledge.

Then there are the real, deep-seated uncertainties. To arrive at any of these projections—average or otherwise—requires taking several steps, each with its own set of known and, most vexingly, unknown unknowns. Uncertainties exist around the amounts of global warming pollutants we emit, the link between emissions and atmospheric concentrations, the link between concentrations and temperatures, the link between temperatures and physical climate damages, the link between physical damages and their consequences, and, at least as important, how society will respond: what coping measures will be undertaken, and how effective they will prove to be.

Nailing down one of these steps—the link between concentrations and eventual temperature increases—has proven particularly elusive. The past three decades of amazing advances in climate science have gotten us no closer to pinpointing the true answer. Double the carbon dioxide concentrations in the atmosphere—something that will surely happen, unless we enact ambitious climate policies now—and eventually global average temperatures are *likely* to go up by between 1.5 and 4.5°C (2.7 and 8°F). Our confidence in that range has increased, but what's now called the "likely" range hasn't changed since the late 1970s, a fact we will revisit in chapter 3, "Fat Tails."

The very term "fat tails" also points to another problem: 1.5 to 4.5°C (2.7 to 8°F) is "likely" in the best sense of that word. The chance is good that we will indeed find ourselves somewhere in that range for how temperatures react when concentrations double, what's known as "climate sensitivity." But there's also a chance we won't. The Intergovernmental Panel on Climate Change (IPCC) describes anything below 1°C (1.8°F) as "extremely unlikely." That assessment is pretty believable, given that the world has already warmed by 0.8°C (1.4°F), and we haven't even yet doubled carbon dioxide concentrations from preindustrial levels. (The 400 ppm that the world just passed is a 40 percent increase over preindustrial levels of 280 ppm.) There's also a chance that final temperatures from a doubling of carbon dioxide concentrations will end up above 4.5°C (8°F). It's "unlikely," but we can't discount the possibility.

Meanwhile, global average warming of 4.5°C (8°F) is beyond the pale of most imagination. Recall the camels in Canada, or at least a planet that none of us would recognize.

But that 4.5°C (8°F) doesn't yet tell the full story. Climate sensitivity describes what happens when concentrations of carbon dioxide in the atmosphere double. What if carbon dioxide concentrations more than double? The International Energy Agency (IEA) predicts levels of 700 ppm, or two-and-a-half times preindustrial levels. Now we are looking at a "likely" range of temperatures between 2 and 6°C (3.6 and 11°F).

Climate science warns that average global warming above 2°C (3.6°F) could trigger potentially devastating events. It's unclear what label to use for global average warming of 6°C (11°F): "catastrophic" no longer seems to do it justice. Mark Lynas, who has painstakingly detailed climate impacts degree by frightening degree, ends his book *Six Degrees* just there. The introduction to the final chapter on 6°C (11°F) begins with a reference to Dante's Sixth Circle of Hell. HELIX, a recently started project funded by the European Union, aims to determine global and regional impacts of specific levels of temperature rise. It, too, ends at 6°C (11°F). And per our own calculations in chapter 3, we are looking at an eventual chance of around 10 percent of *exceeding* that mark.

Whenever science points to the very real potential of these types of catastrophic outcomes, cognitive dissonance kicks in. Facts might be facts, the reasoning goes, but throwing too many of them at you at once will all but guarantee that you will dismiss them out of hand. It just *feels* like it can't or shouldn't be true.

That fickleness of human nature and the limits of our understanding are at the core of the climate policy dilemma. Smarts alone don't seem to make much of a difference here. Solving the dilemma will take a completely different way of thinking.

THE BATHTUB PROBLEM

Think of the atmosphere as a giant bathtub. There's a faucet—emissions from human activity—and a drain—the planet's ability to absorb that pollution. For most of human civilization and hundreds of thousands of years before, the inflow and the outflow were in relative balance. Then humans started burning coal and turned on the faucet far beyond what the drain could handle. The levels of carbon in the atmosphere began to rise to levels last seen in the Pliocene, over three million years ago.

What to do? That's the question John Sterman, an MIT professor, asked two hundred graduate students. More specifically, he asked what to do to stabilize concentrations of carbon dioxide in the atmosphere close to present levels. How far do we need to go in turning off the faucet in order to stabilize concentrations?

Here's what not to do: stabilizing the flow of carbon into the atmosphere today won't stabilize the carbon already there at close to present levels. You're still adding carbon. Just because the inflow remains steady year after year, doesn't mean the amount already in the tub doesn't go up. Inflow and outflow need to be in balance, and that won't happen at current levels of carbon dioxide in the tub (currently at 400 ppm) unless the inflow goes down by a lot.

That seems like an obvious point. It also seems to get lost on the average MIT graduate student, and these students aren't exactly "average." Still, over 80 percent of them in Sterman's study seem to confuse the faucet with the tub. They confuse stabilizing the inflow with stabilizing the level.

To be fair, these two hundred MIT students weren't told about the bathtub analogy. They just saw an excerpt of the "Summary for Policymakers" from the latest IPCC report at the time. That's the document that's meant to explain the issue to our elected officials. If as fundamental a point as the difference between annual emissions and concentrations in the atmosphere—the difference between the inflow and the level of carbon in the tub—is lost on MIT graduate students, what hope is there for the rest of us?

Sure, it's a "Summary for *Policymakers*." Jane Q. Public may not need to understand it, as long as policy makers do. But there, too, is a hiccup. MIT graduate students may well be a good proxy for (better-educated) policy makers. Moreover, there are policy makers, and there are policy makers. The anonymous bureaucrat writing the actual policies may have a Ph.D. in the subject for which he or she is making policy. One hopes. The elected official is unlikely to be a specialist in any particular subject. And ultimately, of course, Jane Q. Voter decides how that person ought to think about a particular issue.

It shouldn't come as a surprise then that one all too popular option among elected officials is a so-called wait-and-see approach to tackling global warming pollution. It's precisely what it sounds like, and it's as misguided as the bathtub analogy would suggest. We can't wait until the moment when that crucial Antarctic ice sheet slips into the ocean and brings us 3 meters (10 feet) closer to where global sea levels were in the Pliocene. At that point, even the last holdouts would realize we are in a climatic emergency. But the emergency is linked to the concentration of carbon in the atmosphere. Society can most directly control the inflow of emissions, and even turning that inflow to zero immediately wouldn't solve the problem. It will take centuries and millennia for the excess carbon to flow out naturally. "Wait and see" might as well be called "give up and fold."

Climate change requires an entirely new way of thinking, something as seemingly foreign to MIT graduate students as to policy makers and the general public. And lest we think getting serious about climate change is as simple as understanding the bathtub analogy and acting accordingly—as seemingly difficult as that alone is—this analogy highlights only two of the Big Four issues: the long-term nature of climate change, with a whole lot of irreversibility mixed in. Nothing yet on the other two: how global and uncertain climate change truly is. The global nature of global warming all but guarantees that deliberately turning off the faucet is incredibly tough to do. Uncertainty doesn't exactly help either, even though it ought to prompt stronger action today. If you don't know precisely how far the tub is from flooding, it's only prudent to turn off the faucet sooner.

WE CAN DO THIS

There are plenty of angles to take from here.

One can try to be optimistic. Yes, things are dire, but look at all the progress. The price of solar panels has declined by 80 percent within five years. Much of it has happened on the backs of German and Chinese households, whose governments took to direct subsidies to bring down costs, but the best way to respond may be to brush up on your German and Chinese for those thank-you notes. They took the hit, for the rest of us to enjoy cheaper solar energy.

Solar energy is not a perfect replacement for fossil sources, at least without significant improvements to electricity market structures and storage technologies. A coal or gas plant can be turned on and off, but we can't control when the sun shines. Still, on a sunny Sunday afternoon, when the sun is up and demand is down, Germany gets 50 percent of its electricity from the sun. Averaged over the entire year of 2013, Germany got almost 5 percent of its electricity from the sun. That's Germany, the industrial powerhouse in Europe, not typically thought of as a particularly sunny place.

Things are looking up globally, too. The world added almost 40 gigawatts (GW) of total solar capacity in 2013, on top of the 30 GW added in 2012, which came on top of the 30 GW added in 2011. The absolute numbers are large, but the rate of change is even more significant. In 2000, the world had around 1 GW of total installed solar capacity. At the end of 2010, the world had 40 GW. By the end of 2013, the tally stood at 140 GW. That's explosive growth on overdrive.

And the all-important policy changes are happening as we speak. None yet is sufficient in itself, but together they provide an impressive array of policy frameworks. Europe has had its carbon market up and (fully) running since 2008. By now, California has the world's most comprehensive carbon market, covering 80 percent of its total greenhouse gas emissions. British Columbia has a carbon tax. China is experimenting with seven regional carbon market trials, and it has a commitment to peak its carbon emissions by 2030. India has a $1-per-ton coal tax. Not a lot, but it's there, and it's positive. Brazil has an ambitious national climate target and has sharply reduced carbon emissions from deforestation. And—since we're being optimistic—in the United States, a solid majority of the electorate would like elected officials to act, at least in principle. A handful more 100-year storms like the two that hit New York City within the course of two years in 2011 and 2012, and we may well see real change.

In fact, the path toward sensible U.S. climate policy is becoming increasingly clear. For one, it will likely go via state capitals like Sacramento. It will also be going through the Clean Air Act and the Environmental Protection Agency's carbon pollution standards for new and existing power plants. At the very least, these regulations could provide a real bargaining chip when it comes to U.S. Congress considering comprehensive climate policy and a direct price on carbon down the line.

||||||||||||||||

Optimism is good. Economics as a discipline is almost pathologically optimistic, even though it's often seen to be a different kind of optimism. Growth is good. Trade is good. Technology is good. There are asterisks for every one of these statements, but they are just that. Few economists may believe that solar panels will save the day, but new technologies have pulled us out of deep environmental morasses in the past—quite literally. New technologies solved the horse manure crisis threatening to engulf New York City at the end of the 19th century. The internal combustion engine banished horses and buggies to taking tourists around Central Park. No one predicted that particular invention at the time. And it didn't require much in terms of active policy intervention: invent car + find oil = *Eureka!*

There may well be one of these breakthroughs just around the corner. Human history seemingly shows that there always is. It's why we are still here as a species. But hoping for a breakthrough is not a strategy. That's why we return to the undeniable importance of policy. That, too, has worked in the past.

For many pollutants, things first got (and are getting) worse, before they got (or will get) better. When Cleveland's Cuyahoga River caught on fire, so did the nascent environmental movement in the United States in the 1960s. This, in turn, led Richard Nixon to sign into law the National Environmental Policy Act of 1969 and create the Environmental Protection Agency. And that was just the beginning. In addition, Nixon went on to sign the Clean Air Act of 1970, the Clean Water Act in 1972, and the Endangered Species Act in 1973, to name just the major ones. A dozen more laws helped round out the "environmental decade." And the U.S. Congress has acted boldly since, with large bipartisan majorities. George H. W. Bush signed into law the Clean Air Act Amendments of 1990. Among others, they led to measures that slashed the pollution that causes acid rain.

All of that applies to local pollutants: the mercury knocking a few points off your kids' IQ, the soot causing them to develop early asthma, the smog making their eyes water and killing their grandparents early, and the toxins in water making it unsafe for anyone to drink. You see, smell, or feel the problem. You petition your government. It reacts. Problem solved.

In reality, it is, of course, much messier than this simple chain would suggest. Niccolò Machiavelli put it succinctly in *The Prince*, published in 1532: "There is nothing more difficult to take in hand, more perilous to conduct, or more uncertain in its success, than to take the lead in the introduction of a new order of things. Because the innovator has for enemies all those who have done well under the old conditions, and lukewarm defenders in those who may do well under the new."

London experienced its first major bout with air pollution in the 1280s. King Edward I established the first air pollution commission in 1285. In 1306, he made it illegal to burn coal. The punishment for repeat offenders: death. You'd think that with the right amount of monitoring and enforcement, this should have taken care of the problem. Alas, the law was soon vacated—and coal-burning has continued ever since.

Never mind all that messiness. Assume for argument's sake that addressing conventional pollutants is as easy as "see something, say something," before watching the rule of law wield its gavel. Climate change just isn't anything like local air pollution. It is, after all, more global, long-term, irreversible, and uncertain than any other environmental problem. The usual politics don't apply. For one, we don't all even agree on the problem. Reverend Martin Luther King Jr. had his dream when the nightmare was clear to most everyone at the time. We don't seem to be quite there yet on the climate front, at least not in the United States.

NO WE CAN'T

Everything we know about the basic chemistry and physics of how our atmosphere works, and everything we know about the economics of how people behave and the messy politics of how we govern ourselves, leads us to believe that things will get worse before they get better. The fact that pumping carbon dioxide into the atmosphere traps heat—the greenhouse effect—had been discovered by 1824, shown in a lab by 1859, and quantified by 1896.

By now, humans have accumulated around 940 *billion* tons of carbon dioxide in the atmosphere, and counting, enough for atmospheric concentrations of carbon dioxide to have busted through the 400 ppm mark. Concentrations are still increasing at a rate of 2 ppm a year, and that annual increase itself is still increasing.

Then there's the biggest problem, and once again a rather unique one: That continued march in the wrong direction is due to seven billion of us, or at least the billion or so high-emitters most responsible for the total number. The responsibility rests with everyone and no one. There's no finger to point. The enemy is us, all of us. The politics are messy. It's often tough to be optimistic.

For every positive piece of climate policy news, there seems to be an opposing negative one. Yes, India has a $1-per-ton coal tax. It also has about $45 billion in annual fossil fuel subsidies. China may have seven regional cap-and-trade trials. It, in turn, subsidizes fossil fuel to the tune of $20 billion annually. The world subsidizes fossil fuels at a rate of over $500 billion per year. That is equivalent to an average worldwide subsidy of some $15 per ton of carbon dioxide emissions, with lower subsidies in most developed economies and much higher per-ton subsidies in oil-rich countries like Venezuela, Saudi Arabia, and Nigeria. Every one of these dollars is a step backward for the climate. Far from moving toward the right incentives, we seem to be guiding markets in exactly the wrong direction.

Another reason we don't always take the optimistic path is that from the economic perspective, it's rather well-trodden. We've known what needs to be done for a long time. For one, stop subsidizing fossil fuels. Now. It will be tough to make the politics work. Just ask Nigerian president Goodluck Jonathan, who stopped fuel subsidies in January 2012 and quickly backtracked, at least partially, after nationwide riots. That still doesn't make the policy prescription any less appropriate economically.

Far beyond just stopping fossil fuel subsidies, the overall policy framework needed for addressing climate change is clear and has been for decades.

THE SOLUTION TO CLIMATE CHANGE

No one is going to win the Nobel Prize in economics for finding the solution to climate change. The economist who came up with it died a decade before the first prize was given out, and the Swedes no longer award their prizes posthumously. Arthur C. Pigou identified the general problem and the solution—what's by now known as "Pigouvian taxes." Each of the 35 billion tons of carbon dioxide emitted *this year* causes at least about $40 worth of damages to the planet, possibly much more. The correct—the only correct—approach is to price each and every ton of carbon according to the damage it causes.

The average American emits about 20 tons per year. That's 20 times $40 or at least about $800 per person per year. But no one is suggesting that every American send in an $800 check at the end of the year. In fact, the entire point is not to. Every one of us ought to face the right incentives each time we turn on the heat or the air conditioner or fill up our tank of gas. At $40 per ton of carbon dioxide, that means about 35 cents per gallon of gasoline.

Pigou's crucial insight was that we ought to see and pay these costs right then and there at the pump. That's the only way to create the right incentives and lead us to incorporate the full cost into our daily decisions—and stop privatizing benefits while socializing costs.

The result of such a price on carbon dioxide will be that we use less coal, oil, and natural gas. We'd pollute less. More specifically, with the correct price we'd be polluting the "optimal" amount. That's not necessarily zero. It's certainly much less than where we are now, with one's weight in pollution going into the atmosphere every day and a half for the average American.

That's the policy solution in a nutshell: put an appropriate price on burning carbon that reflects its true cost to society.

You can get there either through a tax or by creating an explicit market for carbon dioxide emissions: cap overall emissions, allocate allowances to major emitters, and let them trade these allowances to establish a market price for pollution—"cap and trade." In a theoretical vacuum without uncertainty, the two approaches yield the exact same result. Economists love to have epic debates about which is the better approach in practice.

Taxes are simpler, one line of reasoning goes. No, they aren't. Look at the thousands of pages of the U.S. tax code.

Taxes get the price of pollution up. That's what we need. Yes, for now. But cap and trade limits emissions. That's the ultimate point. If emissions go down cheaply, all the better.

Taxes provide price certainty. Maybe, assuming no political tampering. But first off, any cap-and-trade system can be designed with price certainty in mind. It's as straightforward as creating a price floor and provisions to prevent prices from going above a certain level. And more importantly, even without any of these design features, cap-and-trade prices tend to vary just the right way: low prices during a recession, when demand for emissions allowances is low. Higher prices when business investment is strong, all the while ensuring that overall emissions decline in line with the cap.

But if cap-and-trade prices go through the roof, or collapse to zero, the entire system gets discredited. Electricity price spikes may have derailed market deregulation for generations. Sure, but we aren't talking about price spikes here. If anything, we'd expect prices that are much lower than expected because industry tends to have ways to innovate its way to lower compliance costs than previously assumed.

Taxes allow for other measures like Corporate Average Fuel Economy (CAFE) standards to show their effect. Under a cap, these types of overlapping regulations may only shift emissions but not actually reduce them. Fair enough. But that only shows the importance of getting a cap in the first place. With one, fewer of these other measures would be necessary.

That's where the debate stands at the moment, though the final chapter has yet to be written. The latest theoretical insights point to how taxes may allow for easier international coordination. In theory at least, negotiating a uniform tax rate, the proceeds from which are retained by each country, engenders an ever-so-subtle way of countervailing the force of the free-rider problem altogether. If we all agreed on a uniform tax rate per unit of carbon dioxide, then raising the tax would hurt me directly by raising my cost of using carbon-emitting energy but would help me because it makes everyone else cut down on their carbon dioxide emissions as well. By contrast, negotiating caps alone creates the clear incentive for wanting laxer caps. Negotiating a uniform, global tax can achieve something close to the global optimal outcome. That, of course, says nothing yet about the politics, which once again are the biggest hurdle.

For now, just remember that, in theory *and* practice, both taxes and cap-and-trade systems implement Pigou's vision that polluters pay when they are doing the polluting and, hence, will pollute less. We, together with most economists, would be fine with either carbon taxes or caps, done correctly.

Now we can have endless discussions about how to get there in real life. How did the Swedes manage to pass the world's first tax on carbon dioxide in 1991? Why did the French fail in their efforts to enact one in 2009? Why did Europe have the world's first major carbon cap-and-trade system? What's taking the United States so long? And why are we still subsidizing fossil fuels to the tune of $15 per ton of carbon dioxide globally, when the right number should be at least $40 per ton going the other direction?

Plenty of disciplines have useful things to say about each of these questions. Political scientists, psychologists, sociologists, and climate science communicators all have their own variations of the crucial question: If—since—science has been telling us that this is such a grave problem, why hasn't the world acted accordingly?

For one, it's incredibly hard to overcome the huge vested interests fighting against Pigou's and most every economist's vision of the ideal world. Simply saying it ought to be doesn't make it so. Instead of shouting "carbon tax" or "carbon cap," economists ought to work constructively with what we have: second-, third-, and fourth-best solutions (and worse) that create all sorts of inefficiencies, unintended consequences, and other problems, but that roll with the punches of a highly imperfect policy world—and may even remove some existing, imperfect policy barriers at the same time.

Electricity grid reform is a good example. Far from sending proper signals to households and businesses, electricity prices get averaged, subsidized, and artificially stabilized for all sorts of reasons—sending distorted price signals all across the grid. Getting a price on carbon would be great, but grid reform is an essential step toward creating a level playing field for energy efficiency, demand response, and renewable energy. It's also a battle that can and needs to be fought entirely outside the U.S. Congress. It's often up to states to set policies. That alone doesn't mean the policy debate will be any more sensible—especially given how much is at stake for traditional, largely fossil-fueled utilities—but it does mean economists ought to engage much more deeply than the standard Pigouvian line about proper carbon pricing.

Gasoline prices paid at the pump are another realm where this discussion between first-best policies and reality plays out in real time. Most every economist's ideal solution to underpriced pollution from driving is to raise the price of gasoline at the pump. But instead of increasing the federal gasoline tax in the United States from 18.4 cents per gallon, its level since 1993, to something closer to the optimal level, the regulatory instrument of choice has been raising corporate average fuel-economy or CAFE standards for cars and trucks. Tightened CAFE standards were likely the one new rule from President Obama's first term with the single biggest climate impact. Opinions differ on how cost-effective CAFE standards are. What's clear is that raising CAFE standards is possible to do, even though increasing the gas tax is theoretically the first-best policy solution. Once again, it would behoove economists to engage in CAFE policy debates much beyond the level of shouting "gas tax" every chance they get.

We won't engage in either exercise. We won't be repeating the "gas tax," "carbon tax," or "carbon cap" mantras every chance we get. We also won't engage in the messy world of electricity grid reform, CAFE standards, and other policy measures that are very much necessary and also require sensible economic thinking.

TOUGHER THAN ANYTHING THAT HAS COME BEFORE

Instead, we'll go back to basic economics and zero in on two topics that move us far beyond the standard debates. In particular, we'll focus on the economics of uncertainty and geoengineering, two topics that are highly uncomfortable, highly charged, and central to understanding why climate change matters to all of us. They also show clearly why we must act now.

Climate change harbors some deep uncertainties, sometimes to the point of actual ignorance. Why don't climate models predict up to 20 meters (66 feet) of sea-level rise and camels in Canada as a result of carbon dioxide concentrations at levels from three million years ago, when the world experienced both? In short, we don't know. But uncertainty is no excuse for inaction. It's a call to tackle the climate problem while we still can.

This is a hellishly difficult problem to solve. And if the world doesn't solve it, it will hit us with full force in unpleasant and unexpected ways. This is where we'll end up:

with the specter of geoengineering. Everything we know about how humans behave, and how they don't, leads us to believe that—unless political leaders muster the courage to act, decisively and soon—the world will inevitably be facing some painful choices. It may be folly to believe that technology (in the form of geoengineering) can, once again, bail out society and the planet from the worst of planetary emergencies. But that's the world we are moving toward.

Talk of geoengineering, much like uncertainty, isn't very comforting. It shouldn't be. It's certainly not an excuse for inaction on sensible climate policy, just as we shouldn't start smoking because an experimental lung cancer drug treatment showed some promise in a lab. The specter of geoengineering should be a clarion call for action. Decisive, and soon.

We will come back to the economics of uncertainties—fat tails—and geoengineering in due course. First, a quick 411 of the other key economic concepts and the general state of the debate that will guide our journey into the unknown, unknowable, and sometimes just plain scary.