10 The Geoengineered Planet

Martin L. Weitzman

If there is one genuinely natural bridge spanning the chasm between now and a century from now, I think it is climate change. We who are here now can envision only the foundation of this bridge, and we see this foundation but through a glass darkly. Even so, we can make out enough features to sense that something big and possibly ominous might be in the making on the distant horizon.

In this chapter, I speculate on the subject of the future of climate change. In particular, I focus on the nexus between humanity and nature as viewed through the lens of geoengineering. Unless there is some currently unforeseeable revolutionary technological breakthrough making large-scale noncarbon energy generation much less expensive in the future, it is extremely difficult for me to imagine a binding international agreement being reached on significant reductions in carbon emissions. By contrast, the temptation may become very great for some medium-developed nation feeling itself under climate change siege to unilaterally geoengineer itself (and the planet) out of high temperatures by seeding the stratosphere with sunshine-reflecting particles because it is so extraordinarily cheap. Such a combination gives rise to two simultaneous public goods problems, creating what I will call the twin externality dilemma of climate change. This twin externality dilemma, and what it might entail, is the central subject of my futuristic chapter.

Frankly, this futuristic intellectual foray belongs in the realm of science fiction. It is much more speculation than prediction. While based on contemporary ideas, it is not a deterministic consequence of them. My hope is that it is good, or at least decent, science fiction, meaning, for me, that it induces us to think big on the connection between the present and the future while also thinking big on the human and natural conditions we

find ourselves embedded in. I also hope the narrative is at least slightly entertaining as well. So this is going to be a big-think piece, opinionated, speculative, a bit rambling, and largely without endnotes.

By narrowing my focus (here to a story about geoengineering and climate change), I believe I am following the broad spirit of Keynes's own short essay, "Economic Possibilities for Our Grandchildren." In my view, Keynes did not stray especially widely in his thoughts, considering the vast potential scope of the subject. Instead, he focused primarily on the particular theme of what it might mean to work less and enjoy leisure more, which for him was a key issue about the distant future. His relative narrowness was probably a wise strategic decision, given the potential immensity of distant future possibilities and the temptation to wander all over the map in exploring them. Here I try a similar narrowing tactic. I do not claim that climate change is the only important issue that the world will face a century from now. Given the vast miasma of uncertainty surrounding the distant future, and especially surrounding climate change in the distant future, I cannot even be sure that this issue will turn out to be as important as I think it is likely to become. Suffice it to say that I have worked on the economics of climate change, that I think it is likely (if not surely) to be highly important for the future of the planet, and that I have a few rambling thoughts about what it all might mean.

Some scientists and others have argued that the geological epoch in which we now live is deserving of a new name, the *Anthropocene*. The basic idea is that we have arrived at a point in geological history where human beings are a prime mover, or even the dominant driver, of changes in the earth system's natural history, on a scale of what we often attribute to titanic underlying natural forces such as ice ages. Parts of this human-induced movement have been going on for some time. The human biomass of some 7 billion people exceeds the biomass of all other large species of animals combined, and indeed it is larger than many entire phyla of the animal kingdom. Our capture of energy for primary production is on the scale of a very large ecosystem. Our footprint on our human-altered planet's landscape is already huge. In all of natural history, no other single species has come anywhere near to dominating planet Earth as *Homo sapiens* has.

Yet it is the possibility—indeed the near inevitability—of human-induced climate change on an enormous planetary scale that, more than anything else, has motivated the concept of the Anthropocene epoch. I

know from personal experience that it can take people a long time to accept the shocking notion that cumulative human emissions of greenhouse gases could be so powerful that they might profoundly change our climate. It can seem too counterintuitive that humans could be instigating global changes of such a magnitude. I think that this is a big part of the problem of accepting the reality of anthropogenic climate change.

The bare facts of climate change will be familiar to most readers. I review some salient features that seem to me particularly relevant for the purposes of where I want to go with this chapter.

Greenhouse gases warm the planet by trapping heat. The most important greenhouse gas by far is carbon dioxide. By massively burning carbon-based fossil fuels, which have been naturally deposited over the course of hundreds of millions of years, humans are emitting carbon dioxide at a rate that is unprecedented even on geological timescales. We know carbon dioxide levels almost exactly for the last 800,000 years from measuring tiny air bubbles trapped in layers of Antarctic ice cores. We are currently (in 2013) experiencing atmospheric carbon dioxide levels that are some 40 percent higher than the highest levels recorded in the previous 800,000 years. If we continue at roughly this growth rate for another century or so, we will likely attain levels of carbon dioxide between two and three times higher than the highest levels attained over the past tens of millions of years. From strong built-in inertias, temperatures and climate changes lag atmospheric carbon dioxide accumulations by centuries (or even millennia) and linger on for centuries (and even millennia) thereafter, even if emissions were to cease altogether. Where we will actually end up a century or so from now and what will be the subsequent consequences depends on many factors that are highly uncertain, but many of the prospects seem unnerving, especially if, as I fear, nothing much significant is done about carbon dioxide emissions.

In this timing sequence, there is a profound dilemma. By geological standards, these greenhouse gas changes are of unprecedented speed. Being geologically instantaneous, and therefore without precedent, what will happen next is highly uncertain and might well involve low-probability, high-impact catastrophic outcomes. However, by standards of a human lifetime, these centuries-spanning changes are incredibly slow, and therefore the consequences seem extremely remote to us. Humans have no experience and limited patience for dealing with even possibly catastrophic

events that unfold this slowly. How much more sense it makes to deal now with the actual problems of today rather than the hypothetical problems of the distant future. So right from the beginning, a fight against global warming faces an unprecedented uphill battle against human nature to convince people to take seriously uncertain, seemingly hypothetical, events that will unfold over the course of centuries from now.

I think it is fair to say that as of today, the overall impact of climate change on the human condition, here and now, has thus far been very small. Climate change just has not yet had an impact on the average person's life. People are therefore being asked to sacrifice in the present in the name of something that to them is hypothetical, that they cannot detect now, whose exact impact is uncertain, and that will become an actual threat to everyday living only in the remote future, if at all. We just do not have experience in this sort of thing. Maybe civilizations that come after ours will have such experience with distant-future threats, but right now we do not.

I touched on the subject of possibly ruinous climate change. The essence of this particularly nasty problem is that there is no natural bound on how bad things might get as greenhouse gases pile up. It is difficult to draw a line in the sand and say that things can only get this bad and no worse. Climate change is unusual in potentially affecting the entire worldwide portfolio of utility by threatening to drive all of planetary welfare to disastrously low levels in the most extreme scenarios. With climate change, all of our eggs are in one basket and there is no obvious way to diversify this macrorisk. It really is possible to wreck the earth with a high enough concentration of greenhouse gases because of the potentially open-ended catastrophic reach of severe climate change. But what is a "high enough concentration of greenhouse gases" that might trigger ruinous climate change? We do not know. Everything is uncertain. We cannot control climate outcomes by controlling greenhouse gas concentrations. At best we can control only the probabilities of climate outcomes by controlling greenhouse gas concentrations. I do not think the world grasps the enormous magnitude of the uncertainties involved in predicting climate change under high-greenhouse-gas scenarios, a perennial source of irritation for a public suspicious of anything but crisp deterministic answers to "what if?" questions.

The massive structural uncertainties at the heart of climate change can seem overwhelming when they are aggregated together. The core problem

is that if we keep emitting GHGs at current rates for another century or so, we will likely be so far outside the range of ordinary geological and human experience that we are unsure what processes will ultimately be unleashed and what will happen next.

The science and economics of climate change consist of a very long chain of tenuous inferences fraught with big uncertainties in every link. Begin with unknown greenhouse gas emissions. We do not know what we should take as a base-case trajectory for future greenhouse gas emissions. We do not know what the world is going to agree on as policy for limiting emissions. And even if the world agreed on some policy or another, there would remain big uncertainties about how available policy levers, like taxes, tradable permits, and standards, will affect actual greenhouse gas emissions.

Then there are big scientific uncertainties about how greenhouse gas flow emissions accumulate through the carbon cycle into atmospheric greenhouse gas stock concentrations. And even if we accurately knew the trajectory of future stocks of atmospheric greenhouse gases, a lot of uncertainty exists about how and when greenhouse gas stock concentrations translate into global average temperature changes. There are also great uncertainties about how global average temperature changes decompose into specific changes in regional and temporal weather patterns, which are much more unsure than any globally averaged number can capture. Here is about where the economic unknowns start to kick in. How are uncertainties about adaptations to, and mitigations of, climate change damages at a regional level translated into regional utility or welfare changes through an appropriate regional "damages function"? What values should be put on the alteration or destruction of existing ecosystems? How should regional utility changes be aggregated into a single worldwide utility function, and what should its overall properties be? What discount rate should be used to convert everything into expected present discounted values?

The result of this lengthy cascading of big uncertainties is a reduced form of truly stupendous uncertainty about the aggregate welfare impacts of climate change. And what I have enumerated here is only a partial list of all of the uncertainties. With climate change a century or so from now representing an extrapolation so far beyond the realms of past experience, there is a large potential scope for things going very wrong that we cannot now even envision—the notorious unknown unknowns whose role here is

very plausibly nonnegligible. Such black swan events represent outliers that are difficult or even impossible to predict. And because there is no natural bound on how bad things might get with catastrophic climate change, some black swan events might have terrible consequences.

The issue of how to deal with the deep structural uncertainties in climate change would be completely different and immensely simpler if systemic inertias, like the time required for the system to naturally remove extra atmospheric carbon dioxide, were short, as is the case for many airborne pollutants like sulfur dioxide, nitrous oxides, and particulates. Then an important component of an optimal strategy might be along the lines of wait and see. With strong reversibility, an optimal climate change policy would logically involve (among other elements) waiting to learn how far out on the bad probability tail the planet will end up, followed by midcourse corrections if we seem to be headed for a disaster. Alas, the problem of climate change seems bedeviled at almost every turn by significant stock accumulation inertias—in atmospheric carbon dioxide, in the absorption of heat or atmospheric carbon dioxide by the oceans, and in many other relevant physical and biological processes, which are slow to respond to attempts at reversal. So the climate change problem is characterized by irreversibilities and the unsureness of being able to learn by our potential mistakes in enough time to reverse change-underlying conditions significantly.

The final background issue on climate change that I present concerns the massive public goods problem posed by this mother of all externalities. Although it is a global public good of immense overall magnitude, climate change has different and extremely uneven impacts on different regions of the world. As a generalization, the poorer and less developed countries are likely to suffer adverse impacts most acutely. Furthermore, when all is said and done, I personally think it will be expensive—perhaps very expensive—to replace a carbon-based energy technology with a non-carbon-based energy technology on a truly worldwide scale. Getting international agreement (with meaningful verification oversight and compliance penalties) on an effective unified strategy for combating global climate change seems like an overwhelmingly difficult task.

The upshot of this discussion of selected background issues on climate change is that I am pessimistic about the prospects for meaningful timely action on averting bad outcomes. Until people actually see and feel that climate change is adversely affecting their daily lives, I fear that not very much

will be done about it. To what extent people will feel that climate change is adversely affecting their daily lives, at what future time they might feel this, what then might be done about it, and whether it might then be too late (or not) are all questions whose answers are plagued by such uncertainty that they are highly speculative. The one thing that I feel fairly secure in predicting is that hypothetical future threats of climate change will not be enough to create a serious worldwide coordinated attack on the problem until the frightening actual impact (or perhaps actual real imminent threat) of some scary geoevent mobilizes a genuine popular groundswell of bottom-up demand for action.

With this pessimistic backdrop, I next explore some of the issues raised by a geoengineered planet a century or so hence.

In one sense the term *geoengineering* might just reflect aspects of the Anthropocene epoch in which humans have already become a prime mover in the earth's natural history. This aspect of geoengineering is largely inadvertent, the accidental by-product of large-scale human planetary alterations for other aims. The other sense of geoengineering connotes purposeful action. This form of purposeful geoengineering is typically aimed at undoing the deleterious planetary alterations that we have already inadvertently geoengineered. More on this later.

To a large extent, we have already inadvertently geoengineered the planet. After all, the geologically instantaneous increase in concentrations of atmospheric GHGs constitutes a massive by-product of the burning of huge past deposits of fossil fuels being converted from carbon into carbon dioxide. So global warming itself is a kind of geoengineering. And if some of the proposed "solutions" to climate change problems were to make a dent in the burning of fossil fuels, they would of necessity involve converting large areas of the planet into wind farms, solar panel assemblages, carbon capture and storage facilities, and the like, with corresponding environmental impacts of their own.

Even without considering climate change, the agricultural need to feed a world population of some 7 billion people has necessitated a geoengineered landscape on a stupendous global scale. Furthermore, housing and transporting and providing work stations for 7 billion of the world's population has also required geoengineered agglomerations of massive proportions. These vast global-scale engineering works will be significantly expanded even more as China, India, and other developing countries seek ever higher

standards of living and as the population of the world increases yet more in the coming century or so.

The upshot is that geoengineering with a big human footprint has been with us for some time and is likely to increase ferociously in the future, no matter how the climate changes. For sure, climate change introduces a whole new dimension to the idea of geoengineering, especially the idea of purposefully geoengineering the planet to undo the inadvertent geoengineering we are doing with massive greenhouse gas releases. There will be much more on this subject later. But the basic concept of massive human planetary alteration is already in place.

Thus, especially with, but even without, geoengineered climate change, humans and nature are already so intertwined that they are no longer separate entities. This is hardly an original thought on my part, but I think that the issue is likely to come to a head in the next century or so with the probable onset of serious climate change. The primary issue here will not so much concern humans preserving nature, but will be more about humans coevolving wisely with nature. I will have more to speculate on this subject when discussing geoengineering in the context of climate change. But this coevolution issue is sufficiently subtle and has snuck up on us so quietly that I want to make sure the idea is anchored by a specific example that is untainted by climate change.

I could have chosen any number of such examples of past large-scale coevolution of humans with nature, but the one I select here for concrete illustration is the North American tallgrass prairie. This vast ecosystem was almost completely destroyed within a generation when it was discovered how to get at the incredibly rich agricultural soils beneath the Midwest prairies with a steel moldboard plow. The North American tallgrass prairie was then nearly instantaneously plowed under for farmland, from end to end. I go into this example in some detail because I want to use it as a kind of paradigm.

What is today the most productive large farmland area in the world was a mere century and a half ago the largest virgin tallgrass prairie in the world. However, it turns out that the now-vanished tallgrass prairie ecosystem of the North American Midwest was itself a relatively recent phenomenon—less than ten thousand or so years old. If a farmer were forced to abandon rich Iowa farmland today, it would pass through several successional stages, with its climax vegetation reverting not to prairie but to

some form of mature woodland. Prairie, rather than woodland, existed in the North American Midwest in the first place because of recurrent periodic fires. These well-documented large-scale burns were set by Native Americans. They were mostly deliberate and primarily for hunting purposes. Today what is called "prescribed burning" is a critical ingredient of all serious prairie restoration projects.

But what exactly is the prescribed burning trying to restore? It is not attempting to resurrect "nature" in the pure sense of what would have been there without human intervention in the first place. That would be a woodland climax ecosystem. Prairie restoration is trying to bring back an ecosystem that coevolved with human intervention. Is this tallgrass prairie ecosystem "natural" or "artificial"? Does restoring it represent reverence for wilderness or reverence for the past? Does it make a difference?

I think that the paradigm of the tallgrass prairie restoration dilemma may become ratcheted up to a central theme a century or so from now if and when large-scale climate change impacts are being felt. In such a tallgrass prairie paradigm writ large, the issue with climate change will not so much concern humans preserving nature as about humans coevolving wisely with nature. To portray sharply the conceivable magnitude of this distant-future dilemma in the climate change context, I purposely pick an extreme "what if?" example of a low-probability high-impact catastrophic possibility—and what might be the reaction to it a century or so from now.

Human-induced climate change is unusual because there is potentially unlimited downside liability. There is no market in which to short the planet as a hedge against catastrophic damages. To put some bite into this possibility, here is just one example of what might possibly go very wrong in a worst-case scenario.

This example of a potential catastrophe concerns possibly disastrous releases over the long run of bad-feedback components of the carbon cycle that are currently omitted from most general circulation models. The chief concern here is that there may be significant, if highly uncertain, supplementary components that conceptually should be added to the global warming feedbacks that are normally considered on shorter timescales. One omitted component is the possibly powerful self-amplification potential of greenhouse warming that is due to heat-induced releases of sequestered carbon. A vivid example is the huge volume of greenhouse gases currently trapped in tundra permafrost and other boggy soils, mostly as methane, a

particularly potent greenhouse gas. A more remote, but even more dangerously vivid, possibility, which in principle should also be included, is heat-induced releases of the even vaster offshore deposits of methane trapped in the form of clathrates.[1]

There is a very small but unknown positive probability over the long run (centuries to millennia) of having destabilized methane from these offshore clathrate deposits seep into the atmosphere if the temperature of the waters bathing the continental shelves increases just slightly. The amount of methane involved is huge, although it is not precisely known. Most estimates place the carbon-equivalent content of methane hydrate deposits at about the same order of magnitude as all other fossil fuels combined. Over the long run, a methane outgassing-amplifier process could potentially precipitate a disastrous strong, positive feedback warming. Even if the methane from melting permafrost and decomposing clathrates were to be rapidly converted into carbon dioxide, the possible outcomes are still worrisome. This mechanism is one leading suspect in the so-called PETM event of some 55 million years ago when carbon dioxide and temperatures spiked in the geological record over the course of some tens of thousands of years. The increase in carbon dioxide during the PETM is comparable to what might be attained in the course of a century or so from now under the business-as-usual burning of fossil fuels. Average worldwide surface temperatures went up during the PETM by maybe about 5 or so degrees Celsius. If it occurred at all as part of the currently unfolding climate change drama, such a large methane-release event would likely take centuries to materialize because the presumed initiator would be the slow-acting gradual warming of permafrost and ocean waters at the depths of the continental shelves. Thus, while it is a low-probability event that might transpire only centuries from now (if at all), the possibility of a climate meltdown from bad permafrost and clathrate feedbacks is not just the outcome of a mathematical theory; it has some real physical basis. Other examples of an actual physical basis for catastrophic threshold outcomes could be cited. Furthermore, with geologically instantaneous massive increases in greenhouse gases, having no precedent in tens (or even perhaps hundreds) of millions of years, there is the dreaded possibility of really bad black swan climate outcomes that cannot yet be cited because no one has yet thought of them.

Suppose for the sake of argument that such an event like massive methane or carbon dioxide releases with strong feedbacks began in earnest a

century or so from now. In this science fiction scenario, we might then become very scared that we were riding along a trajectory leading to high-temperature increases, accompanied by relatively rapid melting or even collapse of the Greenland and West Antarctica ice sheets (or by altering ocean circulation patterns, or severely altering planet-wide precipitation patterns, and so forth). There could well be many other nasty tipping-point surprises, some of which are black swan events that we cannot now even imagine. What might we then do? In the face of rapidly rising temperatures, we might be tempted to try to deliberately geoengineer the planet as a quick fix, which would be sufficient to restore temperatures to safer levels at least temporarily while we try, this time perhaps seriously, to cut back drastically on greenhouse gas emissions and undertake more permanent, if much slower-acting, measures.

A National Academy of Sciences study defined *geoengineering* as "options that would involve large-scale engineering of our environment in order to combat or counteract the effects of changes in atmospheric chemistry." Similarly, a study of the Royal Society defined it as "the deliberate large-scale manipulation of the planetary environment to counteract anthropogenic climate change." There are several possible forms of geoengineering. But as of now, it seems that there is only one type that would offer a quick fix to the problem of increasing temperatures. This form of geoengineering would create a space sunshade by shooting reflective particles into the stratosphere that block out a small but significant fraction of incoming solar radiation.

Henceforth in this chapter, I abuse terminology by identifying the term *geoengineering* specifically with providing an artificial space sunshade. (In the literature, this is classified under "solar radiation management" as opposed to being classified under "carbon dioxide removal.") I could talk about other forms of geoengineering, but in this brief chapter, I think it is allowable and even useful to focus sharply on this one particular form. So from now on when I discuss geoengineering, I am discussing what I will call a "geoengineered sunshade."

A geoengineered sunshade of particles placed in the stratosphere introduces immense difficulties, dangers, and dilemmas of its own making, which I will touch on shortly. Almost no one is advocating this measure as a first line of defense against climate change. But it might have an important niche role as an emergency fallback component in a complete portfolio

of options to deal with global warming. This may prove to be significant if, for reasons that I have already outlined, very little is done about averting climate change until the effects are visibly and tangibly bearing down on us at least as seriously as, say, a major protracted worldwide recession. Besides, consideration of a particular example in the specific form of a geoengineered sunshade may help us now to envision more concretely a possibly important set of issues that could come online a century or so from now, which, after all, is the ultimate purpose of this book.

The planet itself naturally geoengineers a temporary sunshade every time there is an explosive volcanic eruption anywhere on Earth sufficiently powerful to shoot sulfur dioxide precursors into the stratosphere. The resulting aerosol particles that coalesce around the sulfur dioxide reflect back incoming sunlight, thereby lowering Earth's surface temperatures almost immediately. The effect is rapid, if short lived, because the stratospheric aerosols decompose rapidly (and harmlessly, since there is no acid rain from the sulfur dioxide because it is in the stratosphere). The last time this naturally occurring phenomenon transpired was during the eruption of Mount Pinatubo in 1991, which was estimated to have lowered the average surface temperature of the earth by about 0.5 degrees Celsius during the subsequent year, returning to its baseline temperature thereafter.

Scientists had been well aware of this naturally geoengineered sunshade effect for a long time, along with an awareness that humans could in principle emulate this process if they wished to, even possibly improving it by substituting more effective reflective materials than sulfur-dioxide-centered aerosols. But in 2006, the Nobel Prize–winning chemist Paul Crutzen published a serious proposal that humanity should discuss openly the possible role of an artificially geoengineered sunshade in case the feeble attempts at emissions controls might in the future, as then seemed likely (and to me seems even more likely now), fail to prevent adverse climate events from occurring.[2] Crutzen had been awarded his Nobel Prize for researching the inadvertently geoengineered loss of the ozone layer from man-made chlorofluorocarbons, so he brought along high scientific and moral credibility with his proposal.

Since then, discussion about researching and investigating a geoengineered sunshade has grown enormously. It is an extraordinarily controversial idea. To repeat, almost no serious observer is advocating a geoengineered sunshade as a first line of defense against climate change.

But the fact remains that it is the only measure that can lower worldwide surface temperatures immediately, and therefore it represents the only human response that might quickly ward off the catastrophic impacts of accelerating-temperature trajectories. By comparison, carbon dioxide emissions reductions are extremely slow acting on climate change due to very long inertial lags. Even if it could be so ordained instantaneously, a complete cessation of carbon dioxide emissions would be unlikely to fend off many catastrophes by the time that they appeared. Given the magnitude of the international public goods problem involved and considering its expense and uncertainty, many (including me) reluctantly consider it unlikely that significant worldwide greenhouse gas reductions will begin in earnest until and unless the threat of dangerous climate change becomes tangible and imminent.

The other thing that a geoengineered sunshade has going for it is its unbelievably cheap cost. But is this very low cost a good thing or a bad thing? In fact, its extraordinary costlessness turns the geoengineered sunshade into a public goods nightmare of a magnitude that rivals the climate change problem itself. This twin externality dilemma may actually turn out to haunt the future of the planet.

A geoengineered sunshade has a long list of things going against it. First, it will not alleviate in the slightest any problems associated with an abnormally high concentration of carbon dioxide. Chief among these problems is ocean acidification, which would proceed apace. So the wholesale destruction of entire ocean ecosystems, including wiping out coral reefs, would be unaltered. Indeed, it could be argued that a major oceanic extinction event might even be made more likely if the world were lulled into a false sense of biosecurity from a lower global warming of surface temperatures without corresponding declines in carbon dioxide levels.

The full climatological effects of a geoengineered sunshade are highly uncertain. What little we know of what might happen comes mainly from computer simulations. These are naturally dependent on parameterizations, functional forms, and, last but not least, getting the overall structure right. Detailed prediction of regional weather patterns, such as local precipitation events, tends to be an especially weak aspect of numerical global climate models. So while we can be fairly sure that shooting reflective particles into the stratosphere will lower global surface temperatures on average, we are much more highly unsure about what will happen at

various times and various locations on Earth. For many critics, the law of unintended consequences reigns supreme here. Almost for sure, precipitation patterns would be altered, perhaps greatly altered for the worse. There are also possible threats to the ozone layer. If some crucial structural elements were missing from the models, this could turn out to be yet another instance where human hubris ushered in a catastrophic black swan event. In this way of looking at things, there is a high enough chance that the cure may be worse than the disease to warrant abandoning further thinking about any such enterprise.

Even if it worked perfectly, a geoengineered sunshade is only a temporary solution in the sense that it must be continuously renewed. In the case of sulfur dioxide, the aerosol effect lasts only about a year, and so the stratosphere must be continually reseeded with sulfate particles supplied by rockets, balloons, aircraft, or whatever else. This temporary aspect could be a relatively good thing or a relatively bad thing depending on how it is viewed. On the one hand, the process could be throttled back or even stopped at any time, we hope after serious reduction of greenhouse gas emissions has been attained in the interim (although then it is an open question of how much assurance there might be that the system would return to "normal"). On the other hand, to go in the direction of stratospheric seeding is to put the planet on an addictive and potentially very dangerous drug regimen that may not be easy to stop.

Another argument frequently made against a geoengineered sunshade is that it represents a form of moral hazard problem. Seeding the stratosphere with reflective particles is extraordinarily cheap. An argument frequently made against researching or even entertaining the idea of a geoengineered sunshade is that if the public comprehends just how cheap and easy this approach might be, then the public might easily mistake it for an inexpensive "solution" to the problem of anthropogenic climate change. By this logic, even disseminating such information might well lower the political will to take the more arduous route of seriously reducing greenhouse gas emissions. So maybe it is better not to even bring up discussion of this option in the first place.

My purpose here is not to discuss in much detail the pros and cons of an engineered sunshade approach to the climate change problem. I merely want to convey the most elementary knowledge of the basic underlying issues. This approach has currently received sufficient attention that it has already generated a sizable literature, which can readily be consulted online.

By now, I think that the outlines of a scientific consensus have emerged on the role of an engineered sunshade approach to the climate change problem. Not everyone subscribes to what I am describing here as a consensus view, but enough do that I think it is worthwhile giving a rough outline before proceeding further.

Almost everyone agrees that a geoengineered sunshade is a very scary proposition with enough inherent dangers that it is vastly inferior to the more conventional strategy of cutting back severely on greenhouse gas emissions. Almost everyone also agrees that a geoengineered sunshade is likely to be a vastly less expensive way of keeping down average global temperatures than the more conventional strategy of cutting back severely on greenhouse gas emissions. More controversial is what seems like an emerging consensus that the downside risks of not doing research now (or in the near future) on a geoengineered sunshade outweighs the downside risks of undertaking now (or in the near future) early preliminary research on this option. The main argument for doing research is the one originally put forth by Paul Crutzen in 2006, which at the time it was published lacked broad support and was probably opposed by a majority of the scientific community. The argument is simple: the measures thus far undertaken to curtail greenhouse gas emissions are woefully inadequate and look as if they will be woefully inadequate into the foreseeable future. In such an eventuality, what do we turn to should future temperatures rise sharply, accompanied by what looks like the approaching danger of some form of global catastrophe? A geoengineered sunshade is the only option currently imaginable that is capable of knocking down global average temperatures in a hurry. On balance, all things being considered, is it not better to be prepared by finding out as much as possible about this option well before any temptation arises to employ it? So the argument in favor of doing research now on a geoengineered sunshade is really almost an argument by default.

I now push this argument by default even further by emphasizing that the future of a geoengineered sunshade has a certain inevitable unavoidability built into it. This inevitable unavoidability comes from the second, thus far relatively neglected, public goods aspect of this twin externality of climate change.

In a perceptive and important article, "The Incredible Economics of Geoengineering," Scott Barrett drew attention to how unbelievably cheap the economic costs of putting up and keeping in place a geoengineered sunshade would be.[3] Essentially any determined country with even a medium-sized

economy could, if unopposed, put up a geoengineered sunshade on its own in answer to its own perceived need to lower global temperatures quickly.

This is a true twin externality to the conventional externality of curtailing greenhouse gases. The conventional climate change externality is the mother of all externalities because cutting back on greenhouse gases is so expensive relative to the difficulty of attaining meaningful global agreement (with international verification and compliance penalties) on the public good of minimal climate change. But then a geoengineered sunshade might be called the father of all externalities because knocking down global average temperatures is so cheap that one country can do it unilaterally to fit its own particular perceived needs, thereby imposing a dangerous "public bad" on a multitude of other nations.

Let me now move toward the culmination of my speculations on a geoengineered planet. Of course I am writing science fiction. It may well never come to pass for a variety of reasons.

The realistic side of me says that the world will likely only minimally limit greenhouse gas emissions until it is perceived that some clear and present danger of climate change is an actual threat to the average person. The mother of all externalities is too strong. The costs of a non-carbon-based technology seem high, and the extraordinary degree of international cooperation required is not there, at least not yet. The threat of climate change is just too remote, too hypothetical, too far in the distant future to compete with real problems that are present here and now. To ask people to think and act otherwise, in an international context no less, is to ask a lot of human nature.

The way I think this story will unfold, alas, is that the world will continue to emit a lot of greenhouse gases without much genuine abatement until a crisis of climate change is clearly perceived. We are in something like a Malthusian trap here. I fear that the real-world equilibrating mechanism is that greenhouse gases will pile up in the atmosphere until some clear and present climate danger appears that demands something like immediate action. The mother of all externalities will resist a serious worldwide coordinated attack on the problem until the frightening real impact (or perhaps real imminent threat) of some scary geoevent mobilizes a genuine popular groundswell of bottom-up grassroots demands for immediate action.

Enormous structural uncertainties connect emissions trajectories to the timing of what happens long afterward—including the time when direct

climate threats, however they are defined, first appear. What I am saying is that the uncertainty will be resolved more on the side of the random variable time of arrival being realized given the perceived action threshold threat, rather than the other way around. In other words, define first a climate change of sufficient magnitude for the average person to demand immediate action. Then run a business-as-usual trajectory of climate change. The time when significant action is actually undertaken will be about when the uncertain trajectory attains the minimal perceived impact required for action. So the pessimistic side of me says that we will keep on piling up greenhouse gases until that unknown and uncertain future time when the trajectory of direct climate consequences crosses some immediate-action threshold.

But at just about this point, the father of all externalities kicks in. It is so cheap to unilaterally geoengineer a sunshade that it may prove irresistible for those countries that are especially hard hit by adverse climate change and whose population is demanding immediate relief. My nightmare scenario for a century or so from now is that there may be incredible tension between these two twin externalities of climate change. This could really put the future world in a bad place.

What to do about all this? I think it cannot be emphasized strongly enough that the sooner the world recognizes that a geoengineered sunshade could well turn out to be the father of all externalities, the better prepared we will be for the possible consequences. I do not think that ignoring the second twin externality of climate change is a sensible, or even a feasible, option.

First, there is desperate need for some kind of an overarching international framework to deal with issues of a geoengineered sunshade. At this stage it would be premature to speculate in what final form this might or should emerge, but preliminary meetings to discuss the problem and the issues could and should start taking place soon. The eventual aim is to develop rules and regulations along with a governance structure for determining when and how the international community might conceivably use a geoengineered sunshade. It is far too early to say what these rules and regulations might be, but it is not at all too early to air the preliminary issues that might be included in such a framework.

Simultaneously, we need to find out as much as possible as soon as possible about the science of a geoengineered sunshade, including possibly

running some proof-of-concept field trials on a small scale. The moral hazard argument against researching a geoengineered sunshade is that if the public comprehends just how cheap this approach might be, then it might be seduced by an inexpensive "solution" to the problem of anthropogenic climate change. My own sense is just the opposite: if the public perceives that a geoengineered sunshade is being considered and discussed by governments in the international community, for whatever reason, the shock value is more likely to alert people to just how serious the climate change problem really is. If the father of all externalities (a geoengineered sunshade) is backing us into a corner where we must take its prospects seriously, then maybe we should be expending more effort overcoming the mother of all externalities by negotiating a serious international treaty that hammers out the details of expensive cost sharing to severely limit the world's carbon emissions.

To sum up, this is my own science fiction candidate for the world's biggest problem a century or so hence. The large-scale interbraided tension between "humanity" and "nature," which was a long time in building but whose enormous magnitude first really became apparent with what was labeled the Anthropocene epoch at the beginning of the twenty-first century, will widen and strengthen to epic proportions by the beginning of the twenty-second century. At that time there will no longer be any illusions remaining about humans "preserving" nature; the real issue will concern humans coevolving wisely with their own anthropogenic version of a forever altered nature.

The driving force in this transformation will be the tension between the twin externalities of climate change. The mother of all externalities will prove too strong to pay the price today for an expensive global public good whose payoffs are viewed by the average person as hypothetical and located in an abstract distant future. The more pressing weight of other concerns will prevent significant cutbacks being made in carbon emissions now. This situation of relative inaction will more or less continue until sufficient greenhouse gases pile up in the atmosphere over a long enough time that some clear and present climate danger appears that clearly demands something like immediate action. So the pessimistic prediction is that the current trajectory of more-or-less business as usual will continue until the frightening real impact (or perhaps just the real imminent threat) of some scary geoevent mobilizes a genuinely popular groundswell of bottom-up

grassroots demands for immediate relief. When the average person feels that climate change is as immediate a threat to global welfare as, say, a deep, prolonged recession, then strong action will be seriously contemplated.

The time when such a visible climate threat threshold first appears is highly uncertain, being the realization of some incredibly complicated random variable. By this time, however, it may be too late to take effective remedial action. In any event, the father of all externalities will then make the unilateral imposition of an essentially costless geoengineered sunshade extremely tempting for any country feeling especially hard pressed by what it perceives as an intolerable threat to its own well being.

Of course, a great many currently unforeseeable developments during the next century or so could cause this nightmare express to derail and jump off its tracks. But suppose China and India continue to develop at breakneck speed, and suppose their growth momentum is joined by several other less developed countries that also aspire to taste an advanced lifestyle. Suppose the United States and Japan and several other economically advanced countries continue to drag their feet. Then I am not at all sure we can count on the deployment of a carbon-free breakthrough technology or some other miracle to rescue the situation over the next century or so.

I hope that my gloomy scenario never comes to pass. But is it not prudent to simultaneously lobby for the best and plan for the worst? The twin externality dilemma of climate change means that the international community has a lot of work cut out for it over the next century or so.