

The Use and Misuse of Models for Climate Policy

Robert S. Pindyck*

“Pay no attention to the man behind the curtain!”

— L. Frank Baum, *The Wonderful Wizard of Oz*

Introduction

In a recent article (Pindyck 2013a), I argued that integrated assessment models (IAMs) “have crucial flaws that make them close to useless as tools for policy analysis” (page 860). In fact, I would argue that the problem goes beyond their “crucial flaws”: IAM-based analyses of climate policy create a perception of knowledge and precision that is illusory and can fool policymakers into thinking that the forecasts the models generate have some kind of scientific legitimacy. Despite the fact that IAMs can be misleading as guides for policy, they have been used by the U.S. government to estimate the social cost of carbon (SCC) and evaluate tax and abatement policies.¹

What are the crucial flaws that make IAMs so unsuitable for policy analysis? They are discussed in detail in Pindyck (2013a), but the most important ones can be briefly summarized as follows:

- (1) Certain inputs—functional forms and parameter values—are arbitrary, but they can have huge effects on the results the models produce. One example is the discount rate. There is no consensus among economists as to the “correct” discount rate to use in estimating the SCC, but different rates will yield wildly different estimates of the SCC and the optimal amount of abatement that any IAM generates. Differences in inputs, particularly the discount rate, largely explain why the IAM-based analyses of Nordhaus (2008) and Stern (2007) come to such strikingly different conclusions regarding optimal

*Sloan School of Management, Massachusetts Institute of Technology, 100 Main Street, Cambridge, MA 02142; e-mail: rpindyck@mit.edu.

¹The SCC measures the external cost of burning a ton of carbon and thus informs us as to how large a carbon tax (or equivalent policy) should be imposed to reduce carbon dioxide (CO₂) emissions. For details on how the U.S. government estimates the SCC, see the technical reports from the Interagency Working Group on Social Cost of Carbon (2010, 2013), which used three IAMs to arrive at estimates of the SCC. Also see Greenstone, Kopits, and Wolverton (2013) for an illuminating explanation of the process used by the Interagency Working Group to estimate the SCC.

abatement. Because the modeler has so much freedom in choosing functional forms, parameter values, and other inputs, the model can be used to obtain almost any result one desires, thereby legitimizing what might be little more than a subjective opinion about climate policy.

- (2) We know very little about *climate sensitivity*, i.e., the temperature increase that would eventually result from a doubling of the atmospheric carbon dioxide (CO₂) concentration, but this is a key input to any IAM. The problem is that the physical mechanisms that determine climate sensitivity involve crucial feedback loops and the parameter values that determine the strength (and even the sign) of those feedback loops are largely unknown and likely to remain unknown for the foreseeable future.² In fact, as Freeman, Wagner, and Zeckhauser (2015) have shown, over the past decade our uncertainty over climate sensitivity has actually increased.
- (3) One of the most important parts of an IAM is the *damage function*, i.e., the relationship between an increase in temperature and gross domestic product (GDP; or the growth rate of GDP). When assessing climate sensitivity, we can at least draw on the underlying physical science and argue coherently about the relevant probability distributions. But when it comes to the damage function, we know virtually nothing—there is no theory and are no data that we can draw from.³ As a result, developers of IAMs have little choice but to specify what are essentially arbitrary functional forms and corresponding parameter values.
- (4) IAMs can tell us nothing about “tail risk,” i.e., the likelihood or possible impact of a catastrophic climate outcome, such as a temperature increase above 5 °C, that has a very large impact on GDP. And yet it is the possibility of a climate catastrophe that is (or should be) the main driving force behind a stringent abatement policy.

Although many would agree that IAMs have their flaws, some economists (e.g., Metcalf and Stock 2017 and Weyant 2017 in this symposium on the use of IAMs for climate policy) might argue that my assessment of their usefulness for policy analysis is too harsh and that if used properly, the models could help us formulate and evaluate alternative climate policies. Arguments in support of the development and use of IAMs as policy tools include the following:

- (1) All models have flaws—after all, any model is a simplification of reality—and yet economists build and use models all the time. A related argument is that the complexity of climate change and its economic impact makes it especially important to have some kind of model to account for the dynamic interactions among key variables and to guide our thinking about policy design.

²See, e.g., Roe and Baker (2007). Allen and Frame (2007) argue that our lack of knowledge will not change in the coming years and that climate sensitivity is essentially “unknowable.”

³There is a large and growing literature on the use of weather data to estimate the impact of temperature (and other measures of climate), especially with respect to agriculture. For surveys of this literature, see Auffhammer et al. (2013) and Dell, Jones, and Olken (2014). However, these studies are limited to short time periods and small fluctuations in temperature and other weather variables. They do not, for example, describe what has happened over 20 or 50 years following a 5 °C increase in mean temperature and thus do not allow us to specify and calibrate IAM damage functions.

- (2) Yes, there is uncertainty over climate sensitivity, and we know very little about the damages likely to result from higher temperatures. But can't our uncertainty over climate sensitivity or the "correct" damage function be handled by assigning probability distributions to certain key parameters and then running Monte Carlo simulations?
- (3) We have no alternative. We must develop the best models possible in order to estimate the social cost of carbon and/or evaluate particular policies. In other words, working with even a highly imperfect model is better than having no model at all.
- (4) Finally, if we don't use IAMs, how can we possibly estimate the SCC and evaluate alternative greenhouse gas (GHG) abatement policies? Should we rely instead on expert opinion? And don't experts have some kind of implicit mental models that drive their opinions? If so, isn't it better to make the model explicit?

Put simply, even with their faults, can't IAMs still be useful as a tool to help inform policy? That is the question I address in this article. In doing so, I respond to the arguments in support of IAMs that I have just described. In the next section, I briefly discuss the use and misuse of IAMs and explain how and when these models can indeed be helpful versus how they can be misleading. I then address the question of whether our uncertainty about climate sensitivity and climate impacts can be handled by assigning probability distributions to various parameters and then running Monte Carlo simulations. (I argue that the answer is no.) Next I turn to the issue of scientific honesty and transparency. I will argue that the use of IAMs to estimate the SCC or evaluate alternative policies is problematic because it creates a veneer of scientific legitimacy that is misleading. This is followed by a discussion of three additional problems with the use of IAMs and whether these models are our best option for estimating the SCC and evaluating alternative climate policies. Finally, I address the question of what we can rely on—if not IAMs—to formulate climate policy. I argue that the best we can do is rely on "expert" opinion, perhaps combined with relatively simple, transparent, and easy-to-understand models. After all, the ad hoc equations that go into most IAMs are no more than reflections of the modeler's own "expert" opinion.

The Use and Misuse of Climate Change Models

Economists find models useful because they provide a logically consistent way to organize our thinking about the relationships among variables of interest. They help us understand the implications of those relationships and identify the roles of various functional forms and parameter values. That is what made the early efforts at climate change modeling so valuable. The models developed by Nordhaus (1991) and others more than two decades ago were early attempts to integrate climate science with the economic effects of GHG emissions. These early modeling efforts helped economists understand how GHG emissions accumulate in the atmosphere, how that accumulation can affect global mean temperatures, and how higher temperatures might affect GDP and consumption. By including a social welfare function that values the flow of consumption over time, these models can also be used to illustrate the possible welfare effects of different GHG abatement policies and how those welfare effects depend on various parameters.

In effect, these early IAMs can be viewed as pedagogical devices. Indeed, in his book, *The Climate Casino*, Nordhaus (2013) uses his Dynamic Integrated Climate and Economy (DICE)

model to help explain—at a textbook level—how unrestricted GHG emissions can cause climate change and lead to serious problems in the future. He also utilizes the model to illustrate some of the uncertainties we face when thinking about the climate system and when trying to predict the changes to expect under different policies. The book thereby provides students (and others) with a good introduction to climate change policy.

So far, so good. The problem arises when we take these models so seriously that we use them to try to evaluate alternative policies, come up with an “optimal” (i.e., welfare-maximizing) policy, or estimate the SCC. Yes, economists often build and use models to help elucidate the interconnections among variables, but usually they understand and are clear about the limits of those models. They know that a model can help to tell a story in a logically coherent way, but the model might not be able to provide the numerical details of the story. In other words, the model might not be suitable for forecasting or quantitative policy analysis. This is the case for the various versions of the Nordhaus DICE model, as well as the plethora of IAMs (most of which are much more complex than DICE) that have been developed over the past couple of decades. As I explained in Pindyck (2013a), many of the key relationships and parameter values in these models have no empirical (or even theoretical) grounding and thus the models cannot be used to provide any kind of reliable quantitative policy guidance.⁴

Is there any way around this problem? Can IAMs be salvaged as a tool for policy analysis if we somehow account for our lack of knowledge about key relationships and parameter values? I turn to this question next.

The Treatment of Uncertainty

Although some developers of IAMs understand that there is considerable uncertainty over climate sensitivity and that we don't know what the “correct” damage function is, they think they have a solution to this problem. In particular, they believe that the uncertainty can be handled by assigning probability distributions to certain key parameters and then running Monte Carlo simulations. Unfortunately, this won't help. The problem is that we don't know the correct probability distributions that should be applied to various parameters, and different distributions—even if they all have the same mean and variance—can yield very different results for expected outcomes, and thus for estimates of the SCC.⁵

To make matters worse, we don't even know the correct functional forms for some of the key relationships. This is particularly a problem when it comes to the damage function. The damage function used in the Nordhaus DICE model, for example, is a simple inverse

⁴My criticism applies largely to models that are used for benefit–cost analysis (BCA) and estimation of the SCC. Some climate models are used only to analyze long-term climate targets, such as a 2 °C temperature target or a long-term atmospheric CO₂ concentration limit. Such models do not include a damage function (because the models are not used to evaluate the benefits of a temperature or CO₂ concentration limit), and the discount rate is much less important.

⁵In Pindyck (2013a), I took three different but plausible distributions for temperature change: a gamma distribution, a Frechet distribution (also called a Generalized Extreme Value, Type 2 distribution), and the distribution derived by Roe and Baker (2007). I calibrated all three distributions so they have the same mean and variance, and I demonstrated that they imply very different values for the willingness to pay (WTP) to avoid the temperature change. In Pindyck (2007), I discuss the implications of uncertainty for environmental policy more generally.

quadratic relationship:

$$L(T) = 1/(1 + \pi_1 T + \pi_2 T^2), \quad (1)$$

where T is the anthropomorphic increase in temperature and $L(T)$ is the reduction (i.e., the loss) in GDP and consumption for any value of T . (Thus $\text{GDP} = L(T)\text{GDP}'$, where GDP' is what GDP would be if there were no warming.) But remember that this damage function is made up out of thin air. It isn't based on any economic (or other) theory or any data. Furthermore, even if this inverse quadratic function were somehow the true damage function, there is no theory or data that can tell us the values for the parameters π_1 or π_2 , the correct probability distributions for those parameters, or even the correct means and variances.

To illustrate, suppose we (somehow) chose probability distributions for π_1 and π_2 . A Monte Carlo simulation would then give us the expected loss $L(T)$ for any particular temperature increase T . But suppose that we then come to believe that damages are likely to rise very rapidly as T grows, more rapidly than equation (1) would indicate. This might lead us to conclude that the damage function should have a different functional form (e.g., an inverse cubic rather than quadratic). For example, we might decide that the following damage function is preferred:

$$L(T) = 1/(1 + \pi_1 T + \pi_2 T^3). \quad (2)$$

The Monte Carlo simulation will now give us a very different (and larger) expected loss. Likewise, one might argue that we are using the wrong probability distributions for π_1 and π_2 , or that we have the correct distributions but the wrong means and/or variances. Changing the probability distributions or the means and variances of the distributions will also result in a very different estimate of the expected loss.⁶

The basic problem is that we know as little about the correct probability distributions as we do about the damage function to which they are being applied. What can we possibly learn from assigning arbitrary probability distributions to the parameters of an arbitrary function and running Monte Carlo simulations? I would argue that the answer is nothing. The bottom line here is simple: If we don't understand how A affects B, but we create some kind of model of how A affects B, running Monte Carlo simulations of the model won't make up for our lack of understanding.

The Issue of Scientific Honesty

The argument is sometimes made that we have no choice—that without a model we will end up relying on biased opinions, guesswork, or even worse. Thus we must develop the best models possible and then use them to evaluate alternative policies. In other words, the argument is that working with even a highly imperfect model is better than having no model at all. This might be a valid argument if we were honest and up-front about the limitations of the model. But often we are not.

⁶Roughgarden and Schneider (1999) used expert elicitation to obtain information regarding the IAM damage function. Later I discuss a broader and simpler way to use expert elicitation to estimate the SCC.

The Veneer of Scientific Legitimacy

Models sometimes convey the impression that we know much more than we really do. They create a veneer of scientific legitimacy that can be used to bolster the argument for a particular policy. This is particularly true for IAMs, which tend to be large and complicated and are not always well documented. IAMs are typically made up of many equations; these equations are hard to evaluate individually (especially given that they are often ad hoc and without any clear theoretical or empirical foundation) and even harder to understand in terms of their interactions as a complete system. In effect, the model is just a black box: we put in some assumptions about GHG emissions, climate sensitivity, discount rates, etc., and we get out some results about temperature change, reductions in GDP, etc. And although it is not clear exactly what is going on, since the black box is “scientific,” we are supposed to take those results seriously and use them for policy analysis. A couple of examples might help to clarify this point.

The “Limits to Growth” debate

To understand the problem more fully, let’s go back 40 years or so and revisit the “Limits to Growth” controversy. The argument for “Limits to Growth” was based on a simple sequence of ideas that appeared quite reasonable to some environmentalists at the time: (1) The earth contains finite amounts of oil, coal, copper, iron, and other nonrenewable resources. (2) These resources are important inputs for the production of a large fraction of GDP. (3) Because they are finite, we will eventually run out of these resources. In fact, because of the growth of population and GDP, we are likely to run out very soon. (4) When we run out, the world’s developed economies will contract dramatically, greatly reducing our standard of living and even causing widespread poverty. (5) Therefore we should immediately and substantially reduce our use of natural resources (and slow or stop population growth). Although this will reduce our standard of living now, it will give us time to adapt and will push back (or even avoid) that day of reckoning when our resources run out and we are reduced to abject poverty.

Points (1) and (2) are indisputable. Points (3), (4), and (5), however, ignore basic economics. As reserves of oil, copper, and other resources are depleted, the costs of extraction and therefore the prices of these resources will rise, causing their use to decline. Higher prices also create the incentive to find substitutes. Thus we may never actually run out of these resources, although we will eventually stop using them. Most important, given the incentives created by rising prices and the likelihood of finding substitutes, there is no reason to expect the gradual depletion of natural resources to result in economic decline. Indeed, due partly to technological change and partly to the discovery of new reserves, the real prices of most resources have gone down over the past 40 years, and there is no evidence that reserve depletion has been or is likely to be a drag on economic growth.

Although it made little economic sense, the “Limits to Growth” argument gained considerable traction in the press and in public discourse over environmental policy. This was due in part to a lack of understanding of basic economics on the part of the public (and many politicians). But it was also due to the publication and promotion of some simulation models that gave the “Limits” argument a veneer of scientific legitimacy. The most widely promoted and cited models were those of Forrester (1973) and Meadows et al. (1974). These models were actually quite simple; as Nordhaus (1973, 1992) and others explained, they

essentially boiled down to an elaboration of points (1) to (5) above, with some growth rates and other numbers attached. What seemed to matter, however, was that these models required a computer for their solution and simulation. The fact that some of the underlying relationships in the models were completely ad hoc and made little sense didn't matter—the fact that they were computer models made them “scientific” and inspired a certain degree of trust.

“Technical analysis” of the stock market

Another example of an attempt to create a veneer of scientific legitimacy is the “technical analysis” used by stock market analysts to make buy/sell recommendations for particular stocks and for the market as a whole. Sometimes this involves a formal computer model and sometimes just a chart-based “analysis” of the up-and-down movements of what is essentially a random walk. By dressing up the “analysis” with technical-sounding terms like “resistance levels,” “support points,” and potential or actual “breakouts,” the buy/sell recommendations of technical analysts are given a scientific aura: uninformed (or misinformed) investors are led to believe that these recommendations are based on some kind of “science,” even though countless studies have shown that “technical analysis” is totally uninformative.

Overselling the validity of climate models

I do not mean to equate IAMs with the “Limits to Growth” models of the early 1970s, never mind the models used by those who promote the “technical analysis” of stock prices. Developers of IAMs generally do try to base their models' equations as much as possible on climate science and economic principles, and the models I am aware of have much more content than the “Limits to Growth” models. The problem is that climate science and economic principles are limited in what they can tell us about how to specify and parameterize an IAM's equations, which is why the models cannot tell us much about the design of climate policy. Unfortunately, IAM developers and users have sometimes failed to be clear about the models' inadequacies, thereby overselling the validity of the models. The result is that policymakers who rely on the projections of IAMs, and have little or no understanding of how the models are built and how they work, can be misled.

I believe that we need to be much more honest and up-front about the inherent limitations of IAMs. I doubt that the developers of IAMs have any intention of using them in a misleading way. Nevertheless, overselling their validity and claiming that IAMs can be used to evaluate policies and determine the SCC can end up misleading researchers, policymakers, and the public, even if it is unintentional. If economics is indeed a science, scientific honesty is paramount.

Isn't the Use of IAMs the Best We Can Do?

Suppose our job is to come up with an estimate of the SCC, which will be used as the basis for determining the size of a carbon tax. We know that IAMs are deeply flawed, but aren't they still the best game in town? If we acknowledge these flaws and explain that the projections and SCC estimates that are generated have large standard errors, isn't the use of one or more IAMs better than having no model? Not necessarily. Even putting aside the issue of scientific honesty, there are three additional problems with the use of IAMs.

The Modeler Has Too Much Flexibility

Put simply, it is much too easy to use a model to generate, and thus seemingly validate, the results one wants. Take any one of the three IAMs that were used by the U.S. Interagency Working Group (2010, 2013) to estimate the SCC.⁷ With a judicious choice of parameter values (varying the discount rate is probably sufficient), these models will yield an SCC estimate as low as a few dollars per ton, as high as several hundred dollars per ton, or anything in between. Thus a modeler whose prior beliefs are that a stringent abatement policy is (or is not) needed, can choose a low (or high) discount rate or choose other inputs that will yield the desired results. If there were a clear consensus on the correct values of key parameters, this would not be much of a problem. But (putting it mildly) there is no such consensus.

The Interagency Working Group did not try to determine the “correct” values for the discount rate. Instead, it used middle-of-the-road assumptions about the discount rate (setting it at 3 percent) as well as other parameters and arrived at an estimate of around \$33 per ton for the SCC (recently updated to \$39 per ton). But other well-known studies have not used these middle-of-the-road assumptions and have arrived at very different estimates of the SCC. For example, using a version of his DICE model (one of the three models used by the Interagency Working Group), Nordhaus (2011) obtained an estimate for the SCC of \$11 per ton. On the other hand, using the PAGE model, Stern (2007) found an extremely stringent abatement to be optimal, a result that is consistent with an SCC of more than \$200 per ton. Although the models differed in a variety of ways (e.g., the degree of disaggregation and the choice of damage function), the main reason for their wildly different SCC estimates is that Nordhaus used a relatively high discount rate, while Stern used a relatively low discount rate.

The problem here is that there is no consensus regarding the “correct” discount rate. (The Interagency Working Group simply chose a midrange number—3 percent—that the members of the group could all live with; the group’s reports never claimed that this number was in any sense “correct.”) Because reasonable arguments can be made for a low discount rate or a high rate, the modeler simply has too much flexibility in the choice of discount rate. If the modeler is at all biased towards a more or less stringent abatement policy, he/she can choose a discount rate accordingly. Moreover, although I have focused here on the discount rate, as I discuss in Pindyck (2013b), IAMs have other parameters whose choice can lead to a higher or lower SCC estimate.

One might argue that the sensitivity of the SCC to the discount rate is a problem with the definition of the SCC itself rather than the models used to estimate the SCC. Indeed, the SCC is usually defined (and calculated) as the present value of future reductions in GDP resulting from one additional ton of CO₂ emissions today. Given the long time horizon involved, that present value is clearly very dependent on the discount rate. The problem is that this dependence becomes obscured by the many equations of the model.

⁷The three IAMs were DICE (Dynamic Integrated Climate and Economy), PAGE (Policy Analysis of the Greenhouse Effect), and FUND (Climate Framework for Uncertainty, Distribution, and Negotiation). For descriptions of these models, see Nordhaus (2008), Hope (2006), and Tol (2002), respectively.

The Choice of Model is Largely Irrelevant

Suppose we could take away the flexibility that the modeler has in choosing parameters. Perhaps some government agency tells the developer of each model to use a specific set of parameter values, including the discount rate. Does this solve the problem?

You might say that we first need to decide which model to rely on. But the choice of model illustrates a second problem. Let's go back to the wildly different SCC estimates of Nordhaus (\$11) and Stern (\$200 plus). Which SCC estimate should we take to be the correct one? The answer boils down almost entirely to our beliefs about the discount rate and the magnitude of climate sensitivity. In fact, the choice of model—DICE versus PAGE versus some other IAM—doesn't matter all that much. Yes, for a fixed set of parameter values DICE will give a different SCC estimate than PAGE, but the difference will be small compared to the effect of changing the values for the discount rate and climate sensitivity within any one model.

One might argue that DICE and PAGE also differ in terms of their damage functions. That is true; DICE uses the simple inverse quadratic damage function of equation (1) while PAGE uses a more complex and disaggregated set of damage functions. But these damage functions are typically calibrated to give GDP losses for moderate temperature increases (5 °C or less) that match the “common wisdom,” and thus are very similar. The projected GDP losses for very large temperature increases (6 °C or more) do differ significantly across the models. However, in all of the models, such temperature increases occur with low probability and in the distant future, so unless the discount rate is very low, they contribute very little to the estimates of the SCC.

If one believes that we should use market-based discount rates (i.e., the rates we actually observe in financial markets) and a relatively low value for climate sensitivity, then \$11 might be roughly the right number for the SCC. But if instead one believes (perhaps based on some kind of “ethical” argument regarding intergenerational welfare comparisons) that we should use a very low discount rate, then \$200 or so might be the right number, especially if one also uses a larger value for climate sensitivity. The point here is that there is hardly any need for a model; decide on the discount rate and climate sensitivity and you pretty much have an estimate of the SCC. The model itself is almost a distraction.

Why is the SCC largely determined by the discount rate and climate sensitivity rather than by the specific IAM used in the analysis? The answer is that the impact of GHG emissions on climate is, almost by definition, a function of climate sensitivity. In addition, that impact is a very slow and gradual process. Even with no abatement, most studies indicate that any significant warming will not occur for several decades. This means that while the costs of a GHG abatement policy are incurred starting now, most of the benefits will come in the distant future. If those future benefits are discounted at a market-based rate (say around 5 percent), their present value will be very small and thus the implied SCC will be very small. To get a large SCC we need to discount future benefits at a very low rate (say around 1 percent) or use a very large value for climate sensitivity. So, is the SCC small or large? To answer that we only have to agree on climate sensitivity and the discount rate.⁸ We don't necessarily have to agree on which model to use.

⁸Economists are sharply divided on the discount rate that should be used for the analysis of climate change policy and long-lived public projects such as the construction of dams and bridges. There is a large and growing literature on the discount rates (plural, because some argue that the rate should decline over time) that should be used for very long time horizons. For an overview, see Gollier (2013).

Catastrophic Outcomes

As I stated in the introduction and explained in detail in Pindyck (2013a, 2013b), what really matters for the SCC is the likelihood and possible impact of a catastrophic climate outcome: a much larger than expected temperature increase and/or a much larger than expected reduction in GDP caused by even a moderate temperature increase. IAMs, however, simply cannot account for catastrophic outcomes.⁹ It is easy to see why by looking at equations (1) and (2), which are alternative representations of the damage function.

Equation (1) is the damage function used in the DICE model. The parameters π_1 and π_2 are chosen to be roughly consistent with the conventional wisdom regarding the loss of GDP that is likely to result from an increase in temperature (T) in the range of 1° to 4 °C. That conventional wisdom, which might turn out to be totally wrong, puts the loss for these kinds of temperature increases at a few percent. The problem is that these damage functions tell us nothing about what to expect if the temperature increases are larger, for example, 5°C or more. Thus, given the arbitrary nature of equation (1), putting in $T = 5^\circ$ or $T = 7^\circ$ is a meaningless exercise and will tell us nothing about the damages we should expect if the temperature were indeed to increase this much. Because of its cubic term, equation (2) will yield much higher damage numbers for $T = 5^\circ$ or more. However, equation (2) is just as arbitrary as equation (1), and the damage numbers are just as meaningless.

Developers of IAMs sometimes claim that their models do account for catastrophic outcomes because they include “tipping points” in the damage functions, such that the loss of GDP increases very sharply when temperature reaches some threshold. But as with the rest of the damage function, the specification of the threshold and the extent to which GDP decreases when the threshold is crossed are arbitrary and not based on any theory or empirics, and thus they cannot tell us much about would happen if the temperature increase turns out to be very large. The damage function, with or without “tipping points,” can do little more than reflect the beliefs of the modeler.

How do we know that the possibility of a catastrophic outcome is what really matters for the SCC and the design of climate policy? Because unless we are ready to accept a discount rate that is very small, the “most likely” scenarios for climate change simply don’t generate enough damages—in present value terms—to matter.¹⁰ That is why the Interagency Working Group, which used a 3 percent discount rate, arrived at the rather low SCC estimate of \$33 per ton.

Then What Should We Do?

I have argued that IAMs and related models are deeply flawed as tools for estimating the SCC, and in fact can be misleading. But if we don’t rely on IAMs to estimate the SCC and formulate climate policy, what should we do instead? And how can we account for the possibility of catastrophic outcomes?

⁹Stern (2013) provides a detailed discussion of how IAMs grossly underestimate (or ignore) possible catastrophic outcomes. Kriegler et al. (2009) used expert elicitation to estimate tipping points in the climate system and thus potential catastrophic outcomes.

¹⁰I show this formally in Pindyck (2011, 2012). See Weitzman (2011, 2013) for discussions of how “tail risk” affects the SCC, and also the discount rate that should be used to calculate the SCC.

Catastrophic Outcomes and an Average SCC

Focusing on catastrophic outcomes may somewhat simplify the problem of estimating the SCC because we can ignore the most likely outcomes (and disagreements about what those most likely outcomes are) and consider only extreme outcomes. It is important to be clear about what we mean by “catastrophic outcomes.” It is only economic outcomes that matter, not the causes of those outcomes. In other words, it doesn’t matter whether a large decrease in GDP is the result of a dramatic increase in temperature (but a moderate effect of temperature on output) or a moderate increase in temperature (but a dramatic effect of temperature on output). What we have to worry about is the possibility of a climate-induced decrease in GDP so large as to be considered catastrophic. (Of course, climate change could also cause noneconomic damages, such as greater morbidity and mortality, the extinction of species, and social disruptions. I am assuming—as is typically done in the estimation of the SCC—that these noneconomic damages could all be monetized and included as part of the decrease in GDP.)

With this in mind, imagine starting with some scenario for GHG emissions (e.g., no abatement). We could then consider a plausible range of catastrophic outcomes, as measured by percentage declines in GDP broadly defined. Next, what are plausible probabilities that we can attach to these possible outcomes? Here, “plausible” would mean acceptable to a range of economists and climate scientists. Given these plausible outcomes and probabilities, we could then calculate the present value of the benefits from averting those outcomes or reducing the probabilities of their occurrence. In present value terms, the benefits will depend on the discount rate and perhaps other parameters, but if those benefits are sufficiently large and robust to reasonable ranges for those parameters, it would imply a large value for the SCC and support a stringent abatement policy. Let’s denote this present value of benefits by B .

The second step in estimating the SCC would be to ask how great the reduction in annual CO₂ emissions would have to be to avoid a catastrophic outcome, or greatly reduce the likelihood of it occurring. If we sum these annual reductions over some time horizon (say 50 or more years) and denote the total reduction by ΔE , we can then calculate an average measure of the SCC; given B and ΔE , the average SCC is simply $B/\Delta E$.

It is important to stress that this average SCC is different from the marginal SCC that is more commonly used in the climate change literature, and which was estimated by the Interagency Working Group using IAMs. The marginal SCC is the present value of the benefits (i.e., avoided losses in GDP) resulting from reducing this year’s CO₂ emissions by 1 ton, whereas the average measure is the present value of benefits from a much larger reduction in emissions now and throughout the future.¹¹

If we don’t rely on an IAM, how can we estimate B and ΔE and thereby compute an average SCC? I believe that this is best done using expert opinion.

Use of Expert Opinion

In order to determine plausible outcomes and probabilities, and the emission reductions needed to avert these outcomes, we would need to rely on “expert” opinion. For an economist, this is not a very satisfying option. Economists often build models to avoid relying on subjective

¹¹I discuss the advantages of estimating an average as opposed to marginal SCC in Pindyck (2017a, 2017b).

(expert or otherwise) opinions. But it is important to keep in mind that the inputs to IAMs (equations and parameter values) are already the result of “expert” opinion—in this case, the modeler is the “expert.” And, of course, experts are likely to disagree, particularly when it comes to climate change, where our knowledge is so limited. On the other hand, focusing on the extreme tail (i.e., catastrophic outcomes), and the emission reductions needed to eliminate that tail, may reduce the extent of disagreement and, more importantly, will center the debate on what really matters as the driver of policy. Moreover, compared with agreeing on the details of an IAM, it might be easier for climate scientists and economists to reach a consensus on at least a range of answers to the questions raised earlier regarding the likelihood of catastrophic outcomes.

In effect, we would use expert opinion to determine the inputs to a simple, transparent, and easy-to-understand model (and I stress the importance of easy-to-understand). As an example of how this might be done, start with three or four potential catastrophic outcomes that, under business as usual, might occur, say, 50 years in the future. Those outcomes might be a 10, 30, or 50 percent drop in GDP and consumption (or something worse). Now attach probabilities to those outcomes, say .2, .1, and .05, respectively (so the probability of no catastrophe is .65). Given these outcomes and probabilities, and given a discount rate, we can calculate the present value of the expected benefits from avoiding these outcomes. Next, come up with an estimate (or set of estimates and associated probabilities) of the reduction in CO₂ emissions needed to eliminate the catastrophic scenarios. A simple ratio then gives us an estimate of the SCC. Of course, the result will still depend on the discount rate that is used, so we might use a range of discount rates. However, as shown in Pindyck (2017a, 2017b), the average SCC is much less dependent on the discount rate than the marginal SCC, because in the average SCC both future benefits and future costs (emission reductions) are being discounted.

Yes, the calculations I have just described constitute a “model,” but it is a model that is exceedingly simple and straightforward and involves no pretense that we know the damage function, the feedback parameters that affect climate sensitivity, or other details of the climate–economy system. And yes, some experts might base their opinions on one or more IAMs, on a more limited climate science model, or simply on their research experience and/or general knowledge of climate change and its impact. But that is fine because we would be using a range of expert opinions to summarize our current understanding of catastrophic climate outcomes and the range of disagreement over those outcomes.¹²

Some might argue that the approach I have outlined here is insufficiently precise. But I believe that we have no choice. Building and using elaborate models might allow us to think that we are approaching the climate policy problem more scientifically, but in the end, like the Wizard of Oz, we would only be drawing a curtain around our lack of knowledge.

¹²In Pindyck (2016, 2017), I discuss the use of expert opinion to estimate the SCC in detail and present the results of a survey of economists and climate scientists and the average SCC estimates that those results imply.

Conclusions

I have stressed that as economists, we need to be honest and forthcoming about what we do and do not know about climate change and its impact. Just as financial economists would not (or should not) try to sell “technical analysis” to investors, environmental economists should not claim that IAMs can forecast climate change and its impact or that IAMs can tell us the magnitude of the SCC.

Atmospheric scientists have made great progress in understanding how weather patterns develop and change, but they don’t claim to be able to forecast next month’s weather or when the next hurricane will arrive. There has also been great progress in our understanding of the drivers of climate, how GHG emissions can affect climate, and (to a much lesser extent) how changes in climate can affect GDP and other economic variables. But that progress still does not enable us to build and use IAMs as tools for forecasting and policy analysis, and we would be deluding ourselves if we thought otherwise.

It would certainly be nice if the problems with IAMs simply boiled down to an imprecise knowledge of certain parameters, because then uncertainty could be handled by assigning probability distributions to those parameters and then running Monte Carlo simulations. Unfortunately, not only do we not know the correct probability distributions that should be applied to these parameters, we don’t even know the correct equations to which those parameters apply. Thus the best one can do at this point is to conduct a simple sensitivity analysis on key parameters, which would be more informative and transparent than a Monte Carlo simulation using ad hoc probability distributions.

This does not mean that IAMs are of no use. As I discussed earlier, IAMs can be valuable as analytical and pedagogical devices to help us better understand climate dynamics and climate–economy interactions, as well as some of the uncertainties involved. But it is crucial that we are clear and up-front about the limitations of these models so that they are not misused or oversold to policymakers. Likewise, the limitations of IAMs do not imply that we have to throw up our hands and give up entirely on estimating the SCC and analyzing climate change policy more generally.

I have argued that the problem of estimating the SCC and formulating climate policy may be somewhat simplified by the fact that what matters most is the possibility of a catastrophic climate outcome. How probable is such an outcome (or set of outcomes) and how bad would it (they) be? And by how much would emissions have to be reduced to avoid these outcomes? I have argued that the best we can do at this point is to come up with plausible answers to these questions, most likely by relying at least in part on numbers supplied by climate scientists and environmental economists, that is, utilize expert opinion. This kind of analysis would be simple, transparent, and easy to understand. It might not inspire the kind of awe and sense of scientific legitimacy conveyed by a large-scale IAM, but that is exactly the point. It would draw back the curtain and help us to clarify our beliefs about climate change and its impacts.

References

- Allen, Myles R., and David J. Frame. 2007. Call off the quest. *Science* 318: 582–83.
- Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. Using weather data and climate model output in economic analysis of climate change. Working Paper 19087, National Bureau of Economic Research.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2014. What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature* 52(3): 740–98.
- Forrester, Jay W. 1973. *World Dynamics*. Cambridge, MA: Wright-Allen Press.
- Freeman, Mark C., Gernot Wagner, and Richard J. Zeckhauser. 2015. Climate sensitivity uncertainty: When is good news bad? Working Paper 20900, National Bureau of Economic Research.
- Gollier, Christian. 2013. *Pricing the Planet's Future*. Princeton, NJ: Princeton University Press.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton. 2013. Developing a social cost of carbon for U.S. regulatory analysis: A methodology and interpretation. *Review of Environmental Economics and Policy* 7(1): 23–46.
- Hope, Chris. 2006. The marginal impact of CO₂ from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern. *Integrated Assessment* 6(1): 19–56.
- Interagency Working Group on Social Cost of Carbon. 2010. Social cost of carbon for regulatory impact analysis under Executive Order 12866. February. U.S. government. <http://www.whitehouse.gov/sites/default/files/omb/inforeg/for-agencies/Social-Cost-of-Carbon-for-RIA.pdf>.
- . 2013. Technical support document: Technical update of the social cost of carbon for regulatory impact analysis. May. U.S. government, <https://www.whitehouse.gov/sites/default/files/omb/assets/inforeg/technical-update-social-cost-of-carbon-for-regulator-impact-analysis.pdf>.
- Kriegler, Elmar, Jim W. Halt, Hermann Held, Richard Dawson, and Hans Joachim Schellnhuber. 2009. Imprecise probability assessment of tipping points in the climate system. *Proceedings of the National Academy of Sciences* 106: 5041–46.
- Meadows, Donella H., Dennis L. Meadows, Jorgen Randers, and William W. Behrens III. 1974. *The Limits to Growth: A Report for the Club of Rome's Project on the Predicament of Mankind*. New York: Universe Books.
- Metcalfe, Gilbert E., and James Stock. 2017. Integrated assessment models and the social cost of carbon: a review and assessment of U.S. experience. *Review of Environmental Economics and Policy* 11(1): 80–99.
- Nordhaus, William D. 1973. World dynamics: Measurement without data. *Economic Journal* 83: 1156–83.
- . 1991. To slow or not to slow: The economics of the greenhouse effect. *Economic Journal* 101: 920–37.
- . 1992. Lethal model 2: The limits to growth revisited. *Brookings Papers on Economic Activity* 2: 1–59.
- . 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press.
- . 2011. Estimates of the social cost of carbon: Background and results from the RICE-2011 model. Working Paper 17540, National Bureau of Economic Research.
- . 2013. *The Climate Casino*. New Haven, CT: Yale University Press.
- Pindyck, Robert S. 2007. Uncertainty in environmental economics. *Review of Environmental Economics and Policy* 1: 45–65.
- . 2011. Modeling the impact of warming in climate change economics. In *The Economics of Climate Change*, ed. G. Libecap and R. Steckel, chap. 2. Chicago: University of Chicago Press.
- . 2012. Uncertain outcomes and climate change policy. *Journal of Environmental Economics and Management* 63: 289–303.
- . 2013a. Climate change policy: What do the models tell us? *Journal of Economic Literature* 51(3): 860–72.
- . 2013b. The climate policy dilemma. *Review of Environmental Economics and Policy* 7(2): 219–37.
- . 2016. The social cost of carbon revisited. Working Paper 22807. National Bureau of Economic Research.

- . 2017. Taxes, targets, and the social cost of carbon. *Economica*, forthcoming.
- Roe, Gerard H., and Marcia B. Baker. 2007. Why is climate sensitivity so unpredictable? *Science* 318: 629–32.
- Roughgarden, Tim, and Stephen H. Schneider. 1999. Climate change policy: Quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy* 27: 415–29.
- Stern, Nicholas. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
- . 2013. The structure of economic modeling of the potential impacts of climate change: Grafting gross underestimation of risk onto already narrow science models. *Journal of Economic Literature* 51(3): 838–59.
- Tol, R. 2002. Estimates of the damage costs of climate change, part I: Benchmark estimates. *Environmental and Resource Economics* 22: 47–73.
- Weitzman, Martin L. 2011. Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy* 5(2): 275–92.
- Weitzman, Martin L. 2013. Tail-hedge discounting and the social cost of carbon. *Journal of Economic Literature* 51(3): 873–82.
- Weyant, John. 2017. Some contributions of integrated assessment models of global climate change. *Review of Environmental Economics and Policy* 11(1): 115–137.