Coase Lecture—Taxes, Targets and the Social Cost of Carbon

By Robert S. Pindyck

Massachusetts Institute of Technology

In environmental economics, the marginal external cost of emitting a pollutant determines the optimal abatement policy, which might take the form of an emissions tax. But the marginal external cost is often difficult to estimate. This is especially the case when it comes to climate change; estimates of the social cost of carbon (SCC) range from around $10 per metric ton to well over $200 per metric ton, and there has been little or no movement toward a consensus number. Partly as a result, rather than an SCC-based carbon tax, climate policy has focused on a set of targets that would put limits on temperature increases or atmospheric CO_2 concentrations, which in turn imply targets for emission reductions. Economics, however, can tell us little about whether such targets are socially optimal. I discuss the trade-off between taxes versus targets as the focus of policy, explain why it has been so difficult to estimate a *marginal* SCC, and suggest an approach to estimating an *average* SCC through the use of expert elicitation. I argue that such an approach could serve as the basis for a harmonized carbon tax.

INTRODUCTION

I am pleased to be giving this year's Coase Lecture, in part because the subject of the lecture is so closely connected to what Ronald Coase (1960) taught us about environmental economics. I am referring to the Coase Theorem, which I am sure most of you teach in your undergraduate microeconomics courses.[1] The theorem says that as long as property rights are well specified and parties can bargain costlessly and to their mutual advantage, the problem of externalities will take care of itself. In particular, the resulting outcome will be efficient, regardless of how the property rights are specified. So if John owns a paper mill that spills toxic effluent into a lake owned by Jane, then the two can reach an agreement (perhaps involving a payment from John to Jane, and/or an investment by John in an effluent-reducing technology), and that outcome will be efficient.

Of course after explaining the Coase Theorem and giving an example like the one above, a student will raise her hand and say that it lacks practical relevance. After all, most lakes are not owned by a single person or firm, and pollution usually imposes a cost on a great many people, so bargaining is not possible. Instead, we need the government to step in and impose a tax or emission quota. The student might also mention that there may be significant uncertainty and disagreement about the damage from the pollution, which would make bargaining difficult even if there were well-specified property rights. (I will discuss such uncertainty and its implications shortly.) Finally, the student might say that one of the biggest environmental problems today is climate change, and the greenhouse gas (GHG) emissions (mostly CO_2) that cause it. A great many people and firms all around the world emit GHGs, and no country owns the Earth's atmosphere, so what possible relevance does the Coase Theorem have for climate change policy?

In fact, for climate change the theorem is quite relevant. The climate negotiations that took place in Paris during December 2015 were an example of the Coase Theorem at work: countries bargained over a target for worldwide emission reductions, which in turn required country-by-country reductions, but the total had to conform to the overall target. That target for emission reductions would in principle balance the external cost of emissions to society as a whole with the corresponding economic benefits from emissions.

In effect, these and most other international climate negotiations involved a process of bargaining between polluters in the aggregate and households, that is, those who are (or will be) harmed by the pollution, also in the aggregate. (Note that the set of polluters includes some of the households that are harmed by the emissions of others, but benefit from their own GHG emissions.)

These aggregations create Coase-like property rights. The polluters collectively 'own' the production processes (factories, power plants, cars, etc.) that generate GHG emissions, and households collectively 'own' the atmosphere into which the emissions are flowing. And like textbook examples of the Coase Theorem, the negotiations involved (potential) monetary payoffs from rich polluters (developed countries) to poor households (developing countries, and especially those most vulnerable to climate change).[2]

The 2015 Paris meetings were just a step in the international climate negotiations that have been going on for some two decades, and will probably continue in the decades to come. Those negotiations focused on an overall target for reductions in CO_2 emissions, and from that a set of (non-binding) pledges of actions that, in the aggregate, would come close to meeting that target. These pledges correspond to country-by-country emission reduction targets, that is, by how much each country should reduce its emissions relative to some base year level. But this focus on emission reductions creates myriad problems.

For example, given that reducing emissions is costly, should a poor country have to reduce its emissions by as much as a rich country? Should a country (rich or poor) whose per capita emission levels are already low have to reduce its future emissions by as much as a country with per capital emissions that are currently very high? And what should be the *overall* target for emission reductions, given the uncertainty over the timing and potential magnitude of climate change impacts? As one might expect, there are no consensus answers to these questions, which is one reason why international climate negotiations have had such limited success.

Rather than negotiate over country-by-country emission reductions, we might do better using a more traditional approach to pollution externalities generally preferred by economists: estimate the social cost of the pollutant, and impose a tax on the pollutant roughly equal to that social cost.[3] In this case the key pollutant is CO_2, so we would need to estimate the *social cost of carbon* (SCC).[4] There are two roughly equivalent alternatives to the tax, namely a quota on the total amount of pollutant emitted, and tradeable emission permits, where the total number of permits is selected to yield the same quantity of emissions as under the quota. (The latter alternative is usually more efficient because it shifts abatement to those with lower abatement costs.) In practice, however, determining the limit on emissions is also based on the pollutant's social cost, in this case the SCC, and thus on the equivalent tax.

Determining the SCC is therefore crucial: given a consensus estimate of the SCC, we can determine the appropriate size of a carbon tax, and use that as the basis for climate policy. As discussed in the next section, there are a number of reasons why a tax-based climate policy can be preferable to negotiating a set of country-by-country emission reductions. The problem is that there is no consensus estimate of the SCC, which makes it difficult to agree on just how large a carbon tax is needed.

In Section II, I explain why estimating the SCC has been so difficult, and how as a result international climate negotiations have focused on *targets*—targets for temperature increases, which translate into targets for the atmospheric CO_2 concentration, and then targets for emission reductions. In Section III, I examine the

nature of the SCC in more detail, and distinguish between a *marginal* SCC and an *average* SCC. I also explain why an average SCC is likely to be more useful as a guide for policy. In Section IV, I propose an analytical framework for estimating an average SCC, and in Section V, I explain how that framework can be implemented via the use of 'expert elicitation' to arrive at several basic numbers. The details of that framework, its implementation via a computer-based survey of economists and climate scientists, and the resulting SCC estimates based on about 1000 responses, are described in detail in Pindyck (2016). Here I illustrate the approach using the survey responses of eleven economists. Section VI concludes.

I. Advantages of a Carbon Tax

Let us assume for the moment that we could come up with a consensus estimate of the SCC, and that the estimate is done on a worldwide basis (i.e. based on climate damages for the entire world, as opposed to a single country). From that SCC, we could determine the carbon tax that should be applied to *all* countries. As argued by Weitzman (2014, 2015) and others, it is likely that a 'harmonized' carbon tax of this sort is a superior policy instrument, because it can better facilitate an international climate agreement.

Why would a harmonized carbon tax be preferable to the country-by-country emission reductions that have been the foundation of ongoing climate negotiations? First and foremost, the negotiations would be over a single number—the size of the tax—as opposed to the much more complex problem of negotiating emission reductions for each and every country. It should be much easier for countries with different interests, and different per capita incomes and emission levels, to agree to a single number as opposed to a large set of numbers. With country-by-country emission reductions, each country has the free-rider incentive to minimize its own reductions and maximize the reductions of other countries. Of course, small countries would still have a free-rider incentive to refuse to take part in a carbon tax regime (as Chen and Zeckhauser (2016) emphasize), but as long as most of the larger GHG emitters do take part, the overall objective of the agreement can still be achieved.

Second, it is difficult to monitor each country's compliance with its agreed-on emission reductions, and even more difficult to penalize a country that does not comply. A harmonized carbon tax goes a long way towards solving the monitoring problem; compared to emission levels, it is much easier to observe whether countries are indeed imposing the tax to which they agreed. And how can we penalize countries that do not comply? In his paper on 'Climate Clubs', Nordhaus (2015) has suggested the imposition of trade sanctions against non-participating or non-complying countries as a way of countering the free-rider problem. While this might increase compliance somewhat, it would also risk escalation into a trade war (and involve major modifications to established trade agreements). But once again, as long as the larger GHG emitters join and comply with the tax agreement, the objectives will be largely achieved.

Third, a tax arising out of an international agreement can be politically attractive, making both agreement and compliance more likely. The tax would be collected by the government of each country, and could be spent in whatever way that government wants. Thus it enables a government to raise revenue at a lower political cost. Taxes of any kind are unpopular in much of the world, but in this case politicians can justify the tax burden by saying 'the devil made me do it'. Finally, an agreement over a harmonized carbon tax can be quite flexible; for example, it need not prevent monetary transfers from rich countries to poor ones, or other forms of side payments.[5]

Whether the focus of climate negotiations shifts to a carbon tax, or remains anchored to an agreement over an *equivalent* reduction in total worldwide emissions (which then requires the more difficult agreement over allocating that total reduction across countries), we need a consensus estimate of the SCC in order to come up with the correct tax or emission reduction. For reasons that I will discuss, so far it has been impossible to obtain such an estimate. In fact, despite the vast amount of research on climate change, estimates of the SCC range from as low as $10 per metric ton to over $200 per metric ton.[6] Thus it would be accurate to say that we have almost no agreed-on view as to the magnitude of the SCC.

II. Targets Versus an SCC-based Tax

A rough consensus estimate of the SCC would considerably facilitate climate negotiations. We could use that estimate as a focal point for negotiating a harmonized carbon tax, or with estimates of supply and demand elasticities for fossil fuels, calculate an equivalent reduction in emissions. Over the last decade, our inability to reach such a consensus estimate of the SCC is one reason why international climate negotiations have focused on *intermediate targets*.

As opposed to 'final' targets for emission reductions, these intermediate targets put a limit on the end-of-century temperature increase, which is then translated into limits on the mid- and end-of-century atmospheric CO_2 concentrations, which in turn are translated into required aggregate emission reductions now and in the coming decades. The targeted temperature increase has been generally specified to be 2° C, on the grounds that warming beyond 2° C would take us outside the realm of temperatures ever observed on the planet, and thus could be catastrophic. Recently, the target has been reduced to 1.5° C, although many analyses indicate that even the 2° C limit is probably infeasible given the current atmospheric CO_2 concentration, current emission levels, and plausible assumptions about possible reductions in emissions during the next two decades.

A limit on the end-of-century temperature increase would seem to obviate the need for an SCC estimate, but as discussed below, it simply replaces the SCC with an arbitrary target that has little or no economic justification. Although a temperature increase above 2° C may indeed go beyond anything that we have observed, we know very little about its potential impact, and there is little or no evidence that the impact would be catastrophic. (One factor limiting the impact is that warming would occur slowly, allowing time for adaptation.) The considerable disagreement over the potential impact of a 2° C temperature increase is directly related to disagreement over the size of the SCC. Thus to understand the focus on a temperature target, I must address the question of why economists and climate scientists have been unable to reach a consensus number for the SCC.

The problem of estimating the SCC

Why has it been so difficult to estimate the SCC and thereby determine an optimal GHG abatement policy? One important factor is the very long time horizon involved. Even with no uncertainty, the time horizon makes the present value of future benefits from current abatement extremely sensitive to the choice of discount rate, and there is considerable disagreement over what the 'correct' discount rate should be. And then there are the very large uncertainties, some of which we cannot even characterize. The more important uncertainties pertain to the extent of warming under current and

expected future GHG emissions, as well as the economic impact of any climate change that might occur. The impact of climate change is especially uncertain, in part because of the possibility of adaptation. We simply do not know much about how worse off the world would be if by the end of the century the global mean temperature increased by 2° C or even 5° C. In fact, we may never be able to resolve these uncertainties (at least not over the next 50 years). It may be that the impact of higher temperatures is not just unknown, but also unknowable—what King (2016) refers to as 'radical uncertainty', or extreme Knightian uncertainty.[7]

Despite these problems, there has been a proliferation of integrated assessment models (IAMs), both large and small. These models have become the standard tool for evaluating alternative climate policies and estimating the SCC.[8] But as I have argued elsewhere, these models have crucial flaws that make them unsuitable for policy analysis.[9] Putting aside the discount rate problem, because of the current limitations of climate science, these models simply make assumptions about *climate sensitivity*, that is, the temperature increase that would result from a doubling of the atmospheric CO_2 concentration. The models likewise make assumptions about the *damage function*, that is, the relationship between an increase in temperature and GDP. And the models, which generally focus on most likely outcomes, tell us nothing about *tail risk*, that is, the likelihood and possible impact of a catastrophic climate outcome, and the key driver of the SCC.

The difficulty with the use of IAMs for policy analysis goes beyond their arbitrary parameter assumptions and *ad hoc* damage functions. The greater problem, discussed in detail in Pindyck (2017), is that they create a perception of knowledge and precision that is illusory, and can mislead policymakers into thinking that the forecasts generated by the models have some kind of scientific legitimacy. The models simply cover over the true extent of how little we know. As King (2016) puts it (in a very different context): 'The fundamental point about radical uncertainty is that if we don't know what the future might hold, we don't know, and there is no point pretending otherwise.' The models pretend otherwise.

A temperature target

This brings us back to the idea of a temperature target of 2° C. The problem is that without a damage function (the weakest part of any IAM), there is no reason to think that 2° C is more justified than some other number. Of course if one believed that the true damage function is essentially flat up to 2° C and then jumps dramatically to a level that we would consider catastrophic, then the 2° C target might make sense.[10] But there is no good reason to believe that the damage function looks like that. In fact, damage function calibrations in the more widely used models take the GDP loss from a 2° C temperature increase to be less than 3%, which we could hardly call catastrophic. Although we expect the function to be convex, we know little beyond that.

So why is an essentially arbitrary temperature target the focus of policy? Because it is something that people can agree on, without having to debate the nature of damages (and the extent of adaptation that would likely limit those damages), never mind the discount rate that should be applied to benefits and costs over horizons of at least 100 years. Whether or not an agreed-on end-of-century temperature target can be justified on economic or climate science grounds, it provides a basis for agreement on atmospheric CO_2 concentration targets and thus targets for overall emission reductions.

Is a temperature target of this sort the best that we can do? Given the difficulty of estimating the SCC, should it be abandoned as the foundation for climate policy design?

If the objective is simply to *do something* about climate change, then a temperature target might make sense. In fact, we may have reached a point where simply doing something is not entirely unreasonable.[11] But for an economist, it is not very satisfying. In the remainder of this paper, I will suggest an alternative approach based on estimating an *average* SCC, which I argue can provide a better foundation for policy than the more conventional *marginal* SCC. In the next section I explain the concept of an average SCC, and then in Section IV I describe a framework for estimating it by using expert elicitation to obtain the necessary inputs. This framework is described in more detail in Pindyck (2016), which also presents estimation results.

III. MARGINAL VERSUS AVERAGE SCC

The most common approach to estimating the SCC uses an IAM or related model to simulate time paths for the atmospheric CO_2 concentration (based on an assumed path of CO_2 emissions), the impact of the rising atmospheric CO_2 concentration on temperature (and perhaps other measures of climate change), and the reductions in GDP and consumption that will result from rising temperatures. The idea is to perturb the assumed time path for CO_2 emissions by increasing current emissions by one ton, and then calculate a new (and slightly lower) path for consumption. The SCC is then the present value of the reductions in consumption over time resulting from that additional ton of current emissions (based on some discount rate). Note that the SCC calculated this way represents the *marginal* external cost of emitting an extra ton of CO_2.

This marginal calculation is consistent with the way environmental economists usually measure the social cost of a pollutant. The marginal external cost of emitting a ton of CO_2 can be added to the marginal private cost (namely the price of the fossil fuel containing 0.2727 tons of carbon plus the cost of burning the fuel to create one ton of CO_2) to obtain the total marginal cost, which would be compared to the marginal benefit (presumably all private) of burning the fuel. In a static context, efficiency can be achieved by imposing a carbon tax sufficient to equate the total marginal cost with the marginal benefit. (In a dynamic context, things are more complicated, as discussed below.) Although a marginal SCC is a more familiar measure of the external cost of burning carbon, an *average* SCC can be a more useful guide for policy.

The marginal SCC

There are three important reasons why the calculation of a marginal SCC may be of limited use for policy. First, the marginal SCC will change over time, even if the underlying technology is completely fixed, that is, even if the price of fuel, the cost of burning it and the benefit from burning it all remain fixed. With a fixed technology, the marginal SCC will generally rise over time. To see why, consider an extreme case in which the damage function depends on the atmospheric CO_2 concentration (rather than the change in temperature), and there is no damage until the concentration reaches a critical level, at which point the damage jumps to some large value. Then the SCC will rise over time for the same reason that the competitive (and socially optimal) price of a depletable resource with constant extraction cost will rise over time. Think of the unpolluted atmosphere as a resource that gets depleted as CO_2 emissions accumulate, with no damages from an increased CO_2 concentration until a threshold is reached, at which point the resource has been depleted. More generally (and realistically), if damages are a convex function of the CO_2 concentration, then the SCC will still rise over time.

This latter case is analogous to the price evolution of a depletable resource when the cost of extraction or the cost of discovering new reserves rises as depletion ensues, as in models such as Pindyck (1978) and Swierzbinski and Mendelsohn (1989).[12]

These changes in the marginal SCC imply that an optimal carbon tax would have to change over time, as would an equivalent quota on CO_2 emissions. This limits the use of the calculated SCC as a guide for policy. It is hard to imagine an agreement on an international climate policy that is based on a carbon tax that changes year by year. The problem would be even more complex for a policy based on emission targets; those too would change over time, as would the allocation of those targets across countries. Agreeing on how and when to change those targets and allocations would be extremely difficult.

The second problem is that the marginal SCC is even limited in terms of guidance for current policy. It can tell us what *today's* carbon tax (or equivalent emission quota) should be, but only under the assumption that total emissions, now and in the future, are on an optimal trajectory. Given that the marginal SCC and thus the optimal emission quota will change over time, this is a strong assumption. One way to deal with this is to solve a dynamic optimization problem based on some reasonably simple IAM, such as in Nordhaus (2008). But that raises a greater problem, which is the need for some kind of IAM to calculate a marginal SCC in the first place. As discussed earlier, given their crucial flaws, IAMs are simply not credible as tools for policy analysis. Despite this, IAMs are the only tools currently in use for the estimation of the marginal SCC.

As mentioned earlier, any IAM-based estimate of the marginal SCC will be highly sensitive to the choice of discount rate, and there is no consensus among economists as to the 'correct' discount rate. This extreme sensitivity to the discount rate is the third major problem with the marginal SCC, and there is no simple way around it. The marginal SCC is the present value of the future losses of GDP (or, in some calculations, consumption) resulting from the emission of one additional ton of CO_2 today. Given that these losses will occur over the distant future, the sensitivity to the discount rate is unavoidable.

Before proceeding, note that an equivalent definition of the marginal SCC is the present value of the future *avoided losses* of GDP—i.e. future *benefits*—from emitting *one less* ton of CO_2 today. In what follows, it will be convenient to express the SCC in terms of future benefits from reduced emissions, rather than losses from increased emissions.

The average SCC

The average SCC is the present value of the flow of benefits resulting from a much larger reduction in emissions now and throughout the future, divided by the total amount of the reduction over the same horizon. Unlike the marginal calculation, there are various ways to calculate an average SCC. For example, how large a reduction should we consider, and over what horizon? Given that the marginal calculation is consistent with the way environmental economists usually measure the social cost of a pollutant, why work with an average number?

There are several reasons why an average SCC is better suited to the design of climate policy. First, given a fixed time horizon (which could be unlimited), we would not expect the average SCC to change over time. This does not mean that it *cannot* change; an unexpected innovation that facilitates adaptation to higher temperatures would cause it to fall. But unlike the marginal SCC, which can change substantially from year to year, the average SCC provides relatively long-term guidance for a tax or emission target policy.

Second, the average SCC is much less sensitive to the choice of discount rate. The marginal SCC is the present value of the flow of benefits from a one-ton change in current emissions; an increase in the discount rate reduces that present value, but does nothing to the one-ton change in emissions. The average SCC, on the other hand, is the present value of a flow of benefits relative to the present value of a flow of emission reductions. (How that second present value can be computed will be addressed shortly.) That creates an offsetting effect of a higher discount rate. As I will show with some numerical examples, the sensitivity to the discount rate is reduced considerably.

Finally, the marginal calculation requires the use of an IAM or related model with its many assumptions regarding the damage function, etc., along with its lack of transparency. Calculating a marginal SCC does not lend itself to expert elicitation—the approach that I will use—because experts cannot tell us what will happen if we reduce emissions today (and only today) by one ton. And even if we had confidence in the particular IAM that is used, the calculated SCC will be sensitive to the assumption made regarding the base-case time path for CO_2 emissions used in the simulations.

As mentioned above, there are various ways in which an average SCC can be defined and estimated. The next section summarizes a definition and approach to estimation that is described in more detail in Pindyck (2016).

IV. Defining and Estimating an Average SCC

My approach to estimating the SCC relies on the elicitation of expert opinions regarding: (1) the probabilities of alternative economic outcomes of climate change, but *not* the causes of those outcomes; and (2) the reduction in emissions needed to avoid or limit those outcomes. For example, a possible outcome might be a 20% or greater reduction in GDP. Whether that outcome is the result of a large increase in temperature but a moderate impact of temperature on GDP, or the opposite, is not of concern. What matters is simply the likelihood of the outcome and the amount of abatement needed to avert it. Also, I am particularly concerned with catastrophic outcomes, that is, climate-caused percentage reductions in GDP that are large. The reason is that unless we are ready to accept a discount rate on consumption that is extremely small (e.g. 1%), the 'most likely' scenarios for climate change simply cannot generate enough damages—in present value terms—to matter.[13] The basic framework, discussed in detail in Pindyck (2016), can be summarized as follows.

1. The primary object of analysis is the economic impact of (anthropomorphic) climate change, measured by the reduction in GDP (broadly defined so as to include indirect impacts such as ecosystem destruction and increased rates of morbidity and mortality). I ignore the mechanisms by which ongoing CO_2 emissions can cause climate change and by which climate change can reduce GDP. I care only about the *outcomes* that can result from CO_2 emissions.
2. I want the probabilities of these outcomes. For example, what is the probability that under 'business as usual' (BAU)—i.e. no significant emissions abatement beyond that mandated by current policy—we will experience a climate-induced reduction in GDP 50 years from now of at least 10%? At least 20%? At least 50%? I rely on expert opinion for the answers.
3. What are the emission reductions needed to avert the more extreme outcomes? Starting with an expected growth rate of CO_2 emissions under BAU, by how much would that growth rate have to be reduced to avoid a climate-induced reduction in

GDP 50 years from now of 20% or more? Once again, I rely on expert opinion for answers.

Different experts will arrive at their opinions in different ways. Some might base their opinions on one or more IAMs, others use their studies of climate change and its impact. The methods used by experts to arrive at their opinions is not a variable of interest; what matters is that the experts are selected according to their established expertise (which, as discussed below, is based on highly cited publications).[14]

Analytical framework

I begin with a distribution for the climate-induced percentage reduction in GDP 50 years from now, which I denote by z. For simplicity, assume for now that the impacts could be reductions of 0%, 2%, 5%, 10%, 20% or 50%, and that according to a hypothetical expert, the probabilities are those given in the top panel of Table 1, where F is the cumulative distribution corresponding to the probabilities in the third row.

Let Y_0 denote what GDP will be if there is no climate change impact, and define $\phi = -\ln(1-z)$. Then a climate change outcome z implies that GDP will be $e^{-\phi}Y_0$. I introduce ϕ because I want to fit several probability distributions to 'expert opinion' damage numbers of the sort shown in Table 1. While z is constrained to $0 \leq z \leq 1$, ϕ is unconstrained at the upper end. Thus I can compare the fits of both fat-tailed (e.g. generalized Pareto) and thin-tailed (e.g. gamma) distributions to such damage numbers, and also compare the implications of these different distributions for SCC estimates.

The top panel of Table 1 applies to a specific horizon $T = 50$ years, but we would expect the impact of climate change to begin before T and continue and increase in magnitude after T. Thus we want to allow for percentage reductions in GDP that increase over time but eventually level out at some maximum value. To simplify the dynamics, I assume that z_t varies over time as follows:

$$(1) \qquad z_t = z_m[1 - e^{-\beta t}].$$

Note that z_t starts at 0 and approaches a maximum value of z_m at a rate given by β. We want to calibrate the maximum reduction z_m and the parameter β.

TABLE 1
PROBABILITIES OF CLIMATE IMPACTS FROM A HYPOTHETICAL EXPERT

			Horizon $T = 50$			
% GDP reduction, z	0	0.020	0.050	0.100	0.200	0.500
$\phi = -\ln(1-z)$	0	0.020	0.051	0.105	0.223	0.693
Probability	0.25	0.50	0.10	0.06	0.05	0.04
$1-F(\phi)$	1	0.75	0.25	0.15	0.09	0.04
			Horizon $T = 150$			
% GDP reduction, z	0	0.020	0.050	0.100	0.200	0.500
$\phi = -\ln(1-z)$	0	0.020	0.051	0.105	0.223	0.693
Probability	0	0.22	0.40	0.20	0.10	0.08
$1-F(\phi)$	1	1	0.78	0.38	0.18	0.08

Begin with β. Suppose that we have average numbers for z_t at two different points in time, T_1 and T_2, and denote the averages by $\bar{z}_1$ and $\bar{z}_2$. The bottom panel of Table 1 shows (hypothetical) probabilities of alternative impacts at a longer horizon, $T_2 = 150$ years. The numbers in Table 1 imply that $\bar{z}_1 = 0.051$ and $\bar{z}_2 = 0.105$. Then from equation (1), we have

$$(2) \qquad [1 - e^{-\beta T_2}]/[1 - e^{-\beta T_1}] = \bar{z}_2/\bar{z}_1 = 2.06.$$

The solution of equation (2) is roughly $\beta = 0.01$. I take this parameter as fixed (non-stochastic).

I treat the maximum impact, z_m, as stochastic. Given β, the distribution for z_m follows from a distribution for z_1, which would be derived from a range of expert opinions (for $T_1 = 50$). Given that distribution, from equation (1) we have

$$(3) \qquad \tilde{z}_m = \tilde{z}_1/[1 - e^{-\beta T_1}].$$

Equation (2) will not have a positive solution for β if $\bar{z}_2/\bar{z}_1$ is too large. With $T_1 = 50$ and $T_2 = 150$, $\bar{z}_2/\bar{z}_1 = 2.06$ implies $\beta \approx 0.01$, but if $\bar{z}_2/\bar{z}_1$ were 3 or greater, then the solution to equation (2) would be negative. In that case, I set $\beta = 0.002$, which implies $\tilde{z}_m \approx 10 \times \tilde{z}_1$.

I assume that absent climate change, real GDP and consumption grow at the constant rate g. Benefits of abatement are measured in terms of avoided reductions in GDP. GDP begins at Y_0 and evolves as $(1 - z_t)Y_0 e^{gt} = Y_0 e^{gt - \phi_t}$, so the loss at time t from climate change is $z_t Y_0 e^{gt} = (1 - e^{-\phi_t}) Y_0 e^{gt}$. Thus the distribution for z_1 yields the distribution for climate damages in each period, and is the basis for the benefit portion of the SCC calculation.

To calculate an SCC consistent with a distribution for ϕ_1, we also need the reduction in GHG emissions required to avoid some range of outcomes. For example, using the numbers in Table 1, we could ask how much emissions would have to be reduced to avoid the very worst or two worst scenarios in the top part of the table. We would then measure benefits as the present value of the expected avoided reduction in the flow of GDP. This, of course, requires a discount rate, which I denote by R.

Estimating the SCC

The estimate of the SCC begins with a scenario for the objective of GHG abatement, which I take to be the truncation of the tail of the impact distribution (in the context of Table 1, eliminating outcomes of $z \geq 0.20$). I focus on eliminating the tail of the impact distribution because eliminating *any* future impact of climate change is probably impossible and thus not an informative scenario. Also, the tail of the distribution accounts for most of the expected damages from climate change under BAU, so avoiding catastrophic damages should be the primary objective of climate policy. Let B_0 denote the present value of the resulting expected avoided reduction in the flow of GDP.

I use equations (1) and (3) to calculate the benefit from truncating the impact distribution to eliminate outcomes of $z \geq 0.20$, which corresponds to $\phi \geq 0.223$. Let $\mathbb{E}_0(z_1)$ denote the expectation of z_1 based on the full distribution of outcomes, and let $\mathbb{E}_1(z_1)$ denote the expectation of z_1 over the truncated distribution. Then the benefit from truncating the distribution is

$$
\begin{aligned}
B_0 &= \int_0^\infty [\mathbb{E}_0(z_t) - \mathbb{E}_1(z_t)] Y_0 e^{(g-R)t} \mathrm{d}t \\
&= Y_0 \left[\frac{\mathbb{E}_0(z_1) - \mathbb{E}_1(z_1)}{1 - e^{-\beta T_1}} \right] \int_0^\infty (1 - e^{-\beta t}) e^{(g-R)t} \mathrm{d}t \\
&= \frac{\beta Y_0 [\mathbb{E}_0(z_1) - \mathbb{E}_1(z_1)]}{(R - g)(R + \beta - g)(1 - e^{-\beta T_1})}.
\end{aligned}
$$

(4)

In equation (4), $\beta Y_0 [\mathbb{E}_0(z_1) - \mathbb{E}_1(z_1)]/(1 - e^{-\beta T_1})$ is the instantaneous flow of benefits from truncating the outcome distribution, and dividing by $(R-g)(R+\beta-g)$ yields the present value of this flow.[15] Also, note that truncating the outcome distribution at time T_1 (i.e. the distribution for z_1) also implies truncating the distribution for z_t at every time t.

Next, we need the 'cost' of this abatement scenario in terms of the total amount of required emission reductions (in tons of CO_2), which I denote by ΔE. Suppose that emissions this year are E_0, and under BAU are expected to grow at the rate m_0. Suppose that the expert consensus is that to eliminate these worst outcomes, the growth rate of emissions must be reduced to $m_1 < m_0$. We want the sum of all future emission reductions, ΔE, which we will compare to B_0. But how should we calculate ΔE? We could simply add up the total reduction in emissions from $t = 0$ to some horizon T, but the horizon T would be arbitrary. Also, if we assume that abatement costs are constant, then it is cheaper in present-value terms to abate more in the future than today. We need to take into account that future abatement costs (like future benefits) must be discounted.

To do this, I assume that the real cost per ton abated is constant over time.[16] Then, irrespective of the particular value of that cost, I can discount future required emission reductions at the same rate R used to discount future benefits (as long as $m_0 < R$). Thus I calculate ΔE as the present value of the flow of emissions at the BAU growth rate m_0 less the present value at the reduced growth rate m_1:

(5)
$$
\begin{aligned}
\Delta E &= E_0 \int_0^\infty \left[e^{(m_0 - R)t} - e^{(m_1 - R)t} \right] \mathrm{d}t \\
&= \frac{(m_0 - m_1)E_0}{(R - m_0)(R - m_1)}.
\end{aligned}
$$

Here $(m_0 - m_1)E_0$ is the instantaneous (current) reduction in emissions, and dividing by $(R - m_0)(R - m_1)$ yields the present value of the flow of emission reductions.[17]

The average SCC is the ratio $B_0/\Delta E$. Using equations (4) and (5), we have

(6) $S = \dfrac{\beta Y_0 [\mathbb{E}_0(z_1) - \mathbb{E}_1(z_1)]/(1 - e^{-\beta T_1})}{(m_0 - m_1)E_0} \times \dfrac{(R - m_0)(R - m_1)}{(R - g)(R + \beta - g)}.$

The first fraction on the right-hand side of equation (6) is the instantaneous SCC, that is, the current benefit (in dollars) from truncating the impact distribution divided by the current reduction in emissions (in metric tons) needed to achieve that truncation. This instantaneous SCC is a flow variable, and the second fraction puts this flow in present value terms. Thus S is the present value of the flows of benefits and emission reductions throughout the future.

Example

Here is a simple example based on the numbers in Table 1 and data for 2013 world GDP and GHG emissions. World GHG emissions (CO_2 equivalent) in 2013 were about 33 billion metric tons. The average annual growth rate of emissions from 1990 to 2013 was about 3%, but almost all of that growth was due to increased emissions from Asia, which are likely to slow over the coming decades, even under BAU. Thus I will assume that under BAU, emissions would grow at 2% annually (so $m_0 = 0.02$). World GDP in 2013 was about $Y_0 = \$75$ trillion. I set $g = 0.02$ as the real GDP growth rate, and use a discount rate of $R = 0.04$. The numbers in Table 1 imply that β in equation (1) is about 0.01.

Suppose that by reducing the growth rate of emissions from $m_0 = 0.02$ to $m_1 = -0.02$ we could avoid the two 'catastrophic' outcomes in Table 1, that is, $z = 0.20$ and $z = 0.50$. In the top part of Table 1, $\mathbb{E}_0(z_1) = 0.05$ and $\mathbb{E}_1(z_1) = 0.022$. (The latter is the expected value of z_1 for the truncated distribution.) From equation (4), the benefit for avoiding these outcomes is $B_0 = 42.36 \times Y_0(0.05 - 0.022) = 1.186 \times Y_0 = \$89 \times 10^{12}$. Given 2013 emissions, and the assumptions that $m_0 = 0.02$, $R = 0.04$ and $m_1 = -0.02$, equation (5) gives $\Delta E = 1.10 \times 10^{12}$ metric tons. With these numbers, the SCC is $B_0/\Delta E = \$81$ per metric ton.

Table 2 shows the SCC and its components for discount rates ranging from 0.025 to 0.060. As one would expect, the benefit B_0 declines sharply as R is increased; this is why estimates of the marginal SCC are so sensitive to the discount rate. But ΔE also declines as R is increased, because the value of future emissions is discounted. The net result is that the (average) SCC declines as R is increased, but not so sharply.

Distributions for outcomes

Expert opinions regarding outcome probabilities can be used to fit several probability distributions. Of interest is which distribution provides the best fit to these 'data', and what are the implications for the SCC. Here I examine three distributions for ϕ: the generalized Pareto, the log-normal and the gamma distributions.

The generalized Pareto is a logical candidate in part because it allows for a fat tail:

$$(7) \qquad f(\phi) = k\alpha(\phi + k^{1/\alpha})^{-\alpha-1}, \quad \phi \geq 0.$$

TABLE 2
SENSITIVITY OF SCC TO DISCOUNT RATE

R	B_0	ΔE	SCC
0.025	712×10^{12}	5.87×10^{12}	\$121
0.030	267×10^{12}	2.64×10^{12}	\$101
0.040	89×10^{12}	1.10×10^{12}	\$81
0.060	26.7×10^{12}	0.41×10^{12}	\$65

Notes
B_0 is the benefit from truncating the distribution for z in Table 1 to eliminate outcomes of $z \geq 0.20$. ΔE is the required total reduction in emissions, with the emission growth rate reduced from $m_0 = 0.02$ to $m_1 = -0.02$. SCC = $B_0/\Delta E$. Also, $\beta = 0.01$, $g = 0.02$ and $T_1 = 50$ years.

The value of α determines the 'fatness' of the tail; if $\alpha > n$, then the first n moments exist. I calculate expectations by integrating to a maximum value of ϕ, $\phi_{max} = 4.6$, which corresponds to $z_{rmmax} = 0.99$. Thus $\mathbb{E}_0(z_1) = 1 - \mathbb{E}_0(e^{-\phi_1})$ in equation (4) is calculated as

$$\mathbb{E}_0(z_1) = 1 - \int_0^{\phi_{max}} k\alpha(\phi + k^{1/\alpha})^{-\alpha-1}e^{-\phi}\mathrm{d}\phi.$$

Also, $\mathbb{E}_1(z_1) = 1 - \mathbb{E}_1(e^{-\phi_1})$, the expectation of z_1 when the distribution has been truncated to eliminate outcomes above some critical limit ϕ_c, is calculated as

$$\mathbb{E}_1(z_1) = 1 - \frac{1}{F(\phi_c)} \int_0^{\phi_c} k\alpha(\phi + k^{1/\alpha})^{-\alpha-1}e^{-\phi}\mathrm{d}\phi.$$

I also fit a log-normal distribution and gamma distribution to the set of outcome probabilities elicited from experts. The log-normal distribution, which approaches zero exponentially and is thus intermediate between a fat-tailed and a thin-tailed distribution, is

$$f(\phi) = \frac{1}{\sqrt{2\pi}\sigma\phi} \exp\left[\frac{-(\ln \phi - \mu)^2}{2\sigma^2}\right], \quad \phi \geq 0,$$

and the (thin-tailed) gamma distribution is

$$f(\phi) = \frac{\lambda^r}{\Gamma(r)}\phi^{r-1}e^{-\lambda\phi}, \quad \phi \geq 0,$$

where $\Gamma(r)$ is the gamma function. I will estimate the parameters of each distribution from a least squares fit of the cumulative distribution to the set of expert opinions regarding outcomes and probabilities, and compare how they fit the 'data' using the corrected R^2. At the end of the next section, I illustrate this using survey responses from eleven experts.

V. THE USE OF EXPERT ELICITATION

Estimating the average SCC requires: (i) the expected rate of growth of GHG emissions under BAU, m_0; (ii) probabilities of alternative climate-induced reductions in future GDP under BAU, from which I will fit an impact distribution; (iii) the reduced growth rate of emissions needed to truncate the impact distribution, m_1; (iv) the most likely climate impact under BAU 50 years from now, z_1, and at a later date, z_2, from which I can determine β in equation (2); and (v) the discount rate, R. I obtain this information from a survey of economists and climate scientists with highly cited publications related to climate change and its impact. The details of this survey are explained in Pindyck (2016). Here I summarize the basic approach, and as an example, I present the responses of eleven experts who attended a recent conference on climate change, along with the implications of those responses for the SCC.

For an economist, relying on expert opinion might not seem very satisfying. Economists often build models to avoid relying on subjective (expert or otherwise) opinions. However, the inputs to IAMs are already the result of expert opinion; the modeller is the 'expert'. This is especially true when it comes to climate change impacts, where theory and data provide little guidance, and expert opinions might best incorporate alternative viewpoints. One could argue that the approach that I am using involves a model of sorts, but it is a model with very few moving parts, and is much more transparent than an IAM-based analysis. The transparency is particularly important— my SCC estimate reflects the opinions, however arrived at, of those with expertise in the field.

Identification of experts and the questionnaire

I want the opinions of people with significant research experience and expertise in climate change and its impact. This can include climate scientists and economists who have worked on climate change, as well as individuals whose focus has been on policy design. To selects experts, I used Web of Science (WoS) to identify journal articles, book chapters, reviews and other publications on climate change and its impacts that were published during the last 10 years. I included publications in five WoS research areas: agriculture, business and economics, environmental sciences and ecology, geology, and meteorology and atmospheric sciences. WoS searched publication titles, abstracts and keywords for particular climate-change-related search terms. (See Pindyck (2016) for the list of search terms and other details of the survey.)

These results were narrowed to include only publications in each research area that were among the top 10% of publication citation counts for each publication year. (This mitigates effects of different citation practices across research areas and the higher numbers of citations expected for earlier publication years.) These publications were used to identify authors in each research area. Next, the lists of authors were pared down so that the percentage of authors in each research area matches the percentage of highly cited publications in that area.[18] After eliminating duplicates, this yielded about 8000 authors, who were contacted via email and asked to respond to an online questionnaire. (The questionnaire is shown in the Appendix.) Of those contacted, approximately 1000 responded and answered the survey questions. The analysis of those responses is in Pindyck (2016), but to illustrate this approach, some preliminary results are described below.

The SCC according to eleven experts

As a test, the questionnaire was given to 20 economists, eleven of whom responded. Their answers are summarized in Table 3. Although this sample is small and more homogeneous than the full set of survey responses, it helps to illustrate how one can obtain parameter estimates for different probability distributions for ϕ. Note that the eleven respondents generally agree about the growth rate of emissions under BAU (m_0), as well as the likely impact on GDP 50 years from now ($\bar{z}_1$). But opinions regarding the probabilities of alternative outcomes, and opinions regarding the likely impact in 2150, vary widely.

Figure 1 shows the least squares fit of the gamma, generalized Pareto and log-normal cumulative distribution functions to the 11 responses to Question 3. Of the three distributions, the Pareto has the highest corrected R^2 (0.559), so I use that to calculate the

TABLE 3
RESPONSES FROM ELEVEN EXPERTS

Expert	Q1 (m_0)	Q2 ($\bar{z}_1$)	Q3					Q4 ($\bar{z}_2$)	Q5 (m_1)	Q6 (R)
			$\geq 2\%$	$\geq 5\%$	$\geq 10\%$	$\geq 20\%$	$\geq 50\%$			
1	0.015	0.040	0.600	0.200	0.050	0.010	0.001	0.100	0.000	0.0250
2	0.030	0.060	0.590	0.480	0.350	0.200	0.040	0.330	−0.028	0.0225
3	0.015	0.080	0.900	0.500	0.050	0.010	0.00001	0.330	−0.035	0.0310
4	0.020	0.050	0.800	0.300	0.050	0.020	0.00	0.150	0.000	0.0100
5	0.020	0.030	0.950	0.250	0.060	0.020	0.002	0.150	0.002	0.0250
6	0.008	0.040	0.810	0.380	0.110	0.020	0.00001	0.175	−0.005	0.0229
7	0.020	0.090	0.900	0.850	0.350	0.200	0.100	0.650	0.000	0.0200
8	0.010	0.020	0.400	0.150	0.050	0.020	0.010	0.100	0.010	0.0200
9	0.020	0.060	0.900	0.700	0.400	0.100	0.030	0.150	0.004	0.0250
10	0.008	0.010	0.050	0.010	0.005	0.0005	0.00001	0.050	−0.005	0.0200
11	0.020	0.040	0.600	0.200	0.050	0.020	0.010	0.080	−0.010	0.0400
Av.	0.017	0.047	0.682	0.365	0.139	0.056	0.018	0.206	−0.006	0.0238

Notes
Questionnaire given to 20 economists, 11 responded. Also, $\bar{z}_1$ is the most likely reduction in GDP for 2066, and $\bar{z}_2$ is the most likely reduction for 2150.

SCC. The estimated parameters of the distribution (equation (7)) were $\hat{\alpha} = 36.00$ and $\hat{k} = 3.436 \times 10^{11}$. The large estimated value of α implies a distribution that is quite thin-tailed.

I calculated an SCC using this distribution together with the average expert opinion for BAU growth rate of emissions ($m_0 = 0.017$), the growth rate of emissions needed to eliminate outcomes of $z_1 \geq 0.20$ ($m_1 = -0.006$), and the discount rate ($R = 0.0238$). I also need a value for β, but the average responses for $\bar{z}_1$ and $\bar{z}_2$ (0.047 and 0.206, respectively) imply that $\beta < 0$, so I set $\beta = 0.002$. These numbers, along with the estimated Pareto distribution, yield a value for the SCC of \$101.24 per metric ton. This is substantially higher than the recent \$39 estimate of the SCC from the US Interagency Working Group.
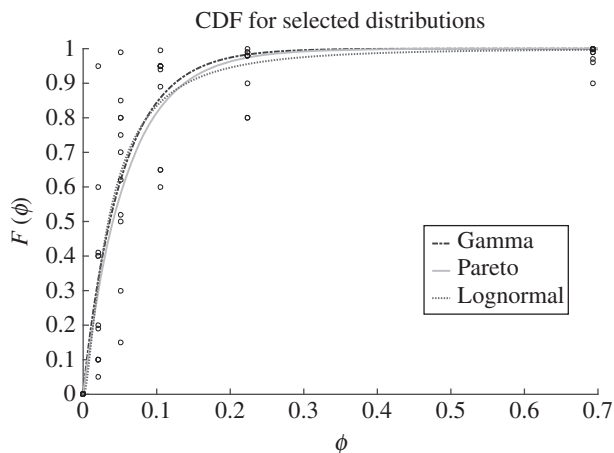

FIGURE 1. Three cumulative distributions, least squares fit to outcome probabilities for eleven experts in Table 3.

VI. Concluding Remarks

For economists, the natural way to think about climate change policy is to determine the external cost of GHG emissions—the social cost of carbon (SCC)—from which an optimal carbon tax can be determined (or tradeable emission permits can be issued based on the total equivalent quota). As I and others have argued, in the context of international climate negotiations, a tax has a number of advantages: it is easier to agree on a single number than a set of country-by-country emission reductions, it is easier to monitor compliance, and it is politically attractive in that each country would retain its own tax revenue. Yet international negotiations have focused on intermediate targets: a maximum temperature increase in 2100 of $2°$ C, from which targets are derived for the maximum atmospheric CO_2 concentration, and in turn total emission reductions (to be allocated across countries). These targets, however, are arbitrary, and agreement on country-by-country emission reductions has been elusive.

Why have we had this focus on targets rather than taxes? In part because despite all the research that has been done, there is no agreement on the magnitude of the *marginal* SCC, which is extremely sensitive to the choice of discount rate and requires an IAM or similar model to estimate. I have argued that as a guide for policy, the marginal SCC is of limited use. It can tell us what *today's* carbon tax should be, assuming that total emissions are on an optimal trajectory, but it will change from year to year. I have introduced an alternative measure, an *average* SCC, which provides a guideline for policy over an extended period of time. I argued that this average SCC can be more useful, especially given the difficult and protracted process for actually agreeing on a climate policy, and it is much less sensitive than the marginal SCC to the choice of discount rate. I proposed an approach to estimating the average SCC that uses expert elicitation to obtain the necessary inputs. Although objections have been raised to the use of expert elicitation, compared to the use of IAMs or related models, it has the advantage of transparency and relative simplicity.

I developed and launched a survey as a way of collecting expert opinion on the inputs to my average SCC calculation, the results of which are presented in Pindyck (2016). My objective, however, is not to obtain a 'final' estimate of the average SCC, but rather to demonstrate how this approach can work and the kinds of answers it can provide. In addition, there are still a variety of problems that remain unresolved. For example, consider the following. (i) What set of possible climate impacts should be presented to survey respondents? More choices, including GDP losses greater than 50%? (ii) Should we fit probability distributions different from the ones I used to the survey responses on impacts? (iii) I used $T_1 = 50$ years and $T_2 = 134$ years (2016 and 2150) as time horizons, but one could argue for alternative horizons. And can experts have meaningful opinions about potential damages as far our as the year 2150? (iv) Are there ways to explicitly include ecosystem destruction, health effects, etc. as part of potential damages?

Also, we must keep in mind that the average SCC is a composite of uncertain parameters (R, m_0, m_1, etc.) and therefore is itself uncertain. And the appeal of the average SCC must be tempered by the fact that it is subject to the same kind of inefficiency that is inherent in average cost pricing for infrastructure. These and other unresolved questions are part of the reason that I view this work as suggestive of an approach, rather than an attempt to arrive at a number that can be used in the next set of climate negotiations. However, unless we are willing to base climate negotiations on a set of arbitrary targets (as is now the case), I see no better alternative to something along the lines of what I have proposed here.

APPENDIX: THE QUESTIONNAIRE

Respondents are asked to read background information and then answer the questions below, skipping those that they cannot or prefer not to answer. They are also asked to indicate on a scale of 1 to 5 the confidence they have in their answers (where 5 is most confident).

- **Introduction**: The purpose of this survey is to estimate the social cost of carbon, an important input to climate policy. Experts, identified from their publications over the past decade, include climate scientists, economists and others who work on climate policy. Respondents' identities will be kept confidential; only overall results of the survey will be published. Before proceeding, read the background information below.
- **Background information**: The questions deal with the impact of climate change and the reductions in GHG emissions needed to limit that impact. 'Impact' and 'emission reductions' should be understood as follows:
    - **Impact**: This is measured as a climate-induced percentage reduction in GDP, broadly defined. Assume that *without* climate change, world real GDP will grow at 2% per year. Climate change, however, could cause floods and other natural disasters, reduce agricultural output, reduce labour productivity, and have other direct effects that would reduce GDPs. Climate change might also have indirect effects, such as ecosystem destruction, social unrest, and increased morbidity and mortality that could further reduce GDP. At issue is *how much lower* future GDP might be as a result of climate change, relative to what it would be without climate change. Is the reduction in GDP likely to be only a few percent, or more than 20% (an outcome some economists would consider 'catastrophic')?
    - **Emission reductions**: While it may be impossible to avoid *any* future impact of climate change, by reducing the growth of GHG emissions we might avoid a very large impact. The average annual growth rate of world GHG emissions over the past 25 years was about 3%, but most of that growth was from Asia. (For the USA and Europe, emissions growth was close to zero.) Some countries have already taken steps to reduce emissions, so under 'business as usual' (BAU)—i.e. if *no additional steps* are taken to reduce emissions—that growth rate might fall to about 2%. However, many experts believe that the growth rate of emissions must drop below this BAU rate to avoid a large impact of climate change. What growth rate of emissions (negative or positive) is needed to avoid a large impact?
- **Question 1**: Under BAU (i.e. no additional steps are taken to reduce emissions), what is your best estimate of the average annual growth rate of world GHG emissions over the next 50 years? (You might believe that the growth rate will change over time; we want your estimate of the *average* growth rate over the next 50 years under BAU.)

  Average emissions growth rate under BAU:

- **Question 2**: If no additional steps are taken to reduce the growth rate of GHG emissions, what is the *most likely* climate-caused reduction in world GDP that we will witness in 50 years? In other words, how much lower (in percentage terms) will world GDP be in 2066 compared to what it would be if there were *no* climate change?

  Most likely percentage reduction in GDP in 2066:

- **Question 3**: Again, suppose that no additional steps are taken to reduce the growth rate of GHG emissions. What is the probability that 50 years from now, climate change will cause a reduction in world GDP of *at least 2%*? (In other words, because of climate change, GDP will be at least 2% lower than it would have been with no climate change.) What is the probability that climate change will cause a reduction in world GDP of at least 5%? At least 10%? At least 20%? At least 50%? Please express each answer as a probability between 0 and 1.

  Probability of 2% or greater reduction in GDP:
  Probability of 5% or greater reduction in GDP:
  Probability of 10% or greater reduction in GDP:

Probability of 20% or greater reduction in GDP:
Probability of 50% or greater reduction in GDP:

- **Question 4**: Now think about the far distant future—the middle of the next century. If no additional steps are taken to reduce the growth rate of GHG emissions, what is the *most likely* climate-caused reduction in world GDP that we will witness in *the year 2150*? In other words, how much lower (in percentage terms) will world GDP be in 2150 compared to what it would be if there were *no* climate change?

Most likely percentage reduction in GDP in 2150:

- **Question 5**: Return to the 50-year horizon, and the possibility that under BAU climate change will cause a reduction in GDP of at least 20%. In Question 1, we asked for your best estimate of the average annual growth rate of GHG emissions over the next 50 years under BAU. What is the average annual growth rate of GHG emissions that would be needed to prevent a climate-induced reduction of world GDP of 20% or more? (By 'prevent', we mean reduce the probability to near zero.) This value might be a positive number, corresponding to slowed growth of emissions, or a negative number corresponding to annual reductions in emissions.

Average emissions growth rate to prevent 20% or greater reduction in GDP:

- **Question 6**: What discount rate should be used to evaluate future costs and benefits from GHG abatement? (Please provide a *single* discount rate.)

Discount rate:

- **Question 7**: Is your expertise primarily in climate science (e.g. how GHG emissions affect climate), primarily in economics (e.g. how climate change can directly or indirectly affect the economy, costs of abatement, policy design, etc.), or in both?

Expertise primarily in climate science, primarily in economics, or both:

ACKNOWLEDGMENTS

NOTES

1. The Coase Theorem is discussed extensively in Chapter 18 of Pindyck and Rubinfeld (2013), and appears in every other microeconomics textbook that I have seen. For a recent discussion of the relevance and application of the Coase Theorem to environmental policy, see Libecap (2016).
2. Of course, the majority of people who would be harmed by additional emissions today are yet to be born, which is perhaps a major difference with the situation that Coase envisaged.
3. Some microeconomics textbooks, e.g. Pindyck and Rubinfeld (2013), define the social cost of an activity as the *total* private plus external cost. In the climate change literature, however, the term social cost usually refers to the external cost alone, so I will use that definition here.
4. The SCC is usually expressed in terms of dollars per ton of CO_2. A ton of CO_2 contains 0.2727 tons of carbon, so an SCC of $10 per ton of CO_2 is equivalent to $36.67 per ton of carbon. The SCC numbers that I present here are always in terms of dollars per ton of CO_2.
5. See Weitzman (2014) for a detailed discussion of these and other aspects of a harmonized carbon tax. Also, see Kotchen (2016) for a discussion of the use of a worldwide SCC versus a domestic (national) SCC.
6. As examples of these extremes, Nordhaus (2011) has estimated the SCC to be $12 per ton of CO_2, so optimal abatement should initially be quite limited. Stern (2007), on the other hand, has estimated the SCC to be about $100 per ton (in today's dollars), but his conclusion that an immediate and drastic cut in emissions is called for is consistent with an SCC above $200.
7. For explanations of why 'radical uncertainty' is likely to apply to climate change, see, for example, Allen and Frame (2007), and Roe and Baker (2007).

8. The US Government's Interagency Working Group (IWG) has used three IAMs to estimate the SCC; see Interagency Working Group on Social Cost of Carbon (2013). For an illuminating discussion of the Working Group's methodology, the models it used, and the assumptions regarding parameters, GHG emissions, and other inputs, see Greenstone *et al.* (2013).

9. For a detailed discussion of these flaws, see Pindyck (2013a,b, (2017).

10. But then what happens after 2100? A global mean temperature that has risen to 2° C by 2100 might be expected to keep rising beyond 2100, so the 2° C limit for 2100 would be too high.

11. Litterman (2013) and Pindyck (2013c) have argued that given the difficulty of reaching a consensus on the SCC, we should simply impose a modest carbon tax, the exact size of which is not very important. This would at least make it clear to politicians and the public that there is indeed a positive external cost of burning carbon that must be added to the private cost. Later, the tax could be adjusted as our understanding of the SCC improves.

12. Still more generally, if technologies that facilitate adaptation arrive stochastically, so that the damage function shifts down, then the marginal SCC can fall. Or if adaptation becomes more difficult and limited than previously anticipated, then the marginal SCC can rise. This is analogous to a depletable resource with a stochastically fluctuating demand curve, so there is a monotonically increasing long-run price trajectory, but the price will fluctuate stochastically around that trajectory. The analogy between the SCC and the price of a depletable resource has been developed in some detail by Becker *et al.* (2011).

13. I have shown this in Pindyck (2012). It is the reason why the Interagency Working Group, which used a 3% discount rate, obtained the low estimate of $33 per ton for the SCC (recently updated to $39).

14. I am not the first to utilize expert opinion as an input to climate policy; see, for example, Kriegler *et al.* (2009), Zickfeld *et al.* (2010) and Morgan (2014). See Oppenheimer *et al.* (2016) for the use of expert opinion to quantify climate uncertainty. For related expert elicitations of the long-run discount rate, see Drupp *et al.* (2015), Weitzman (2001), and Freeman and Groom (2015).

15. In the top panel of Table 1, $\mathbb{E}_0(z_1) = 0.05$. Suppose that by reducing emissions we can eliminate outcomes of $z \geq 0.20$. Increasing the other probabilities so they sum to 1 yields $\mathbb{E}_1(z_1) = 0.022$. Setting $\beta = 0.01$, $g = 0.02$ and $R = 0.04$, this implies $B_0 = 0.00071 Y_0/0.0006 = 1.19 Y_0$. Note that in the *first year*, the benefit from this abatement policy would be less than 0.1% of GDP, but the annual benefit rises over time (as z_t rises), so B_0, the present value of the flow of annual benefits, is greater than current GDP.

16. The real cost per ton abated will be affected over time by two factors that work in opposite directions. Technological progress, e.g. the development of cheaper and better alternatives to fossil fuels, will reduce the cost over time. On the other hand, abatement becomes more and more difficult (and costly) as emissions are continually reduced. It is unclear which of these effects will dominate, so it is reasonable for purposes of estimating the SCC to assume that the cost is constant.

17. Suppose that the growth rate of emissions is reduced from $m_0 = 0.02$ to $m_1 = -0.02$. If $R = 0.04$, then $\Delta E = 0.04 E_0/0.0012 = 33.3 E_0$, i.e. this year's abatement is 4% of current annual emissions, but the present value of all current and future emission reductions is about 30 times this year's emissions.

18. This is done because in some fields (e.g. geology) the authors listed on a paper might include everyone connected with the research, while in other fields (e.g. economics) only primary contributors are included. Thus I identify the research area with the smallest number of authors per publication, and pare down the list of authors in the other areas to match this number, retaining those authors with the most citations.

REFERENCES

ALLEN, M. R. and FRAME, D. J. (2007). Call off the quest. *Science*, **318**, 582–3.

BECKER, G. S., MURPHY, K. M. and TOPEL, R. H. (2011). On the economics of climate policy. *B. E. Journal of Economic Analysis & Policy*, **10**, 1–25.

CHEN, C. and ZECKHAUSER, R. J. (2016). Collective action in an asymmetric world. NBER Working Paper no. 22240.

COASE, R. H. (1960). The problem of social cost. *Journal of Law & Economics*, **3**, 1–44.

DRUPP, M. A., FREEMAN, M. C., GROOM, B. and NESJE, F. (2015). Discounting disentangled: an expert survey on the determinants of the long-term social discount rate. Grantham Research Institute on Climate Change and the Environment Working Paper no. 172.

FREEMAN, M. C. and GROOM, B. (2015). Positively gamma discounting: combining the opinions of experts on the social discount rate. *Economic Journal*, **125**, 1015–24.

GREENSTONE, M., KOPITS, E. and WOLVERTON, A. (2013). Developing a social cost of carbon for use in U.S. regulatory analysis: a methodology and interpretation. *Review of Environmental Economics and Policy*, **7**(1), 23–46.

INTERAGENCY WORKING GROUP ON SOCIAL COST OF CARBON (2013). Technical Support Document: Technical update of the social cost of carbon for regulatory impact analysis. US Government.

KING, M. (2016). *The End of Alchemy: Money, Banking, and the Future of the Global Economy*. New York: W. W. Norton & Co.

KOTCHEN, M. J. (2016). Which social cost of carbon? A theoretical perspective. NBER Working Paper no. 22246.

KRIEGLER, E., HALL, J. W., HELD, H., DAWSON, R. and SCHELLNHUBER, H. J. (2009). Imprecise probability assessment of tipping points in the climate system. *Proceedings of the National Academy of Sciences*, **106**(13), 5041–6.

LIBECAP, G. D. (2016). Coasean bargaining to address environmental externalities. NBER Working Paper no. 21903.

LITTERMAN, B. (2013). What is the right price for carbon emissions? *Regulation*, **36**(2), 38–43.

MORGAN, M. G. (2014). Use (and abuse) of expert elicitation in support of decision making for public policy. *Proceedings of the National Academy of Sciences*, **111**(20), 7176–84.

NORDHAUS, W. D. (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. NBER Working Paper no. 17540.

——— (2015). Climate Clubs: overcoming free-riding in international climate policy. *American Economic Review*, **105**(4), 1339–70.

OPPENHEIMER, M., LITTLE, C. M. and COOKE, R. M. (2016). Expert judgement and uncertainty quantification for climate change. *Nature Climate Change*, **6**, 445–51.

PINDYCK, R. S. (1978). The optimal exploration and production of nonrenewable resources. *Journal of Political Economy*, **86**(5), 841–61.

——— (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.

——— (2013c). Pricing carbon when we don't know the right price. *Regulation*, **36**(2), 43–6.

——— (2016). The social cost of carbon revisited. NBER Working Paper no. 22807.

——— (2017). The use and misuse of models for climate policy. *Review of Environmental Economics and Policy*, **11**(1).

——— and RUBINFELD, D. L. (2013). *Microeconomics*, 8th edn. Boston, MA: Pearson.

ROE, G. H. and BAKER, M. B. (2007). Why is climate sensitivity so unpredictable? *Science*, **318**, 629–32.

STERN, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.

SWIERZBINSKI, J. E. and MENDELSOHN, R. (1989). Exploration and exhaustible resources: the microfoundations of aggregate models. *International Economic Review*, **30**(1), 175–86.

WEITZMAN, M. L. (2001). Gamma discounting. *American Economic Review*, **91**, 260–71.

——— (2014). Can negotiating a uniform carbon price help to internalize the global warming externality? *Journal of the Association of Environmental and Resource Economists*, **1**(1), 29–49.

——— (2015). Internalizing the climate externality: can a uniform price commitment help? *Economics of Energy & Environmental Policy*, **4**(2), 37–50.

ZICKFELD, K., MORGAN, M. G., FRAME, D. J. and KEITH, D. W. (2010). Expert judgments about transient climate response to alternative future trajectories of radiative forcing. *Proceedings of the National Academy of Sciences* **107**(28), 12451–6.