

UNCERTAINTY IN ENVIRONMENTAL ECONOMICS*

by

Robert S. Pindyck

Massachusetts Institute of Technology
Cambridge, MA 02142

This draft: November 2006

Abstract: In a world of certainty, the design of environmental policy is relatively straightforward, and boils down to maximizing the present value of the flow of social benefits minus costs. But the real world is one of considerable uncertainty – over the physical and ecological impact of pollution, over the economic costs and benefits of reducing it, and over the discount rates that should be used to compute present values. The implications of uncertainty are complicated by the fact that most environmental policy problems involve highly nonlinear damage functions, important irreversibilities, and long time horizons. Correctly incorporating uncertainty in policy design is therefore one of the more interesting and important research areas in environmental economics. This paper offers no easy formulas or solutions for treating uncertainty – to my knowledge, none exist. Instead, I try to clarify the ways in which various kinds of uncertainties will affect optimal policy design, and summarize what we know and don't know about the problem.

JEL Classification Numbers: Q28, L51, D81.

Keywords: Environmental policy, uncertainty, irreversibilities, sunk costs, cost-benefit analysis, discount rates, global warming, catastrophic outcomes.

*This paper was written for the *Review of Environmental Economics and Policy*. My thanks to Paul Joskow, Charles Kolstad, Suzanne Leonard, Rob Stavins, Martin Weitzman, and an anonymous referee for helpful comments and suggestions.

1. Introduction

An introductory course in environmental economics typically teaches students that the design and evaluation of a policy to deal with an environmental problem boils down to cost-benefit analysis. The instructor might proceed as follows. Left to their own, humans (i.e., producers and consumers) do bad things to the environment, such as polluting rivers and lakes, spewing sulfur dioxide into the air, and releasing ozone-depleting chlorofluorocarbons (CFCs). Government intervention – restrictions or taxes on emissions, banning CFCs – prevents some of this destructive behavior, and thereby reduces the amount of environmental damage. But it does so at a cost (e.g., electric power producers must install expensive scrubbers, and air conditioners must be made with a more expensive or less efficient refrigerant). So the policy problem boils down to deciding whether the benefit in terms of less environmental damage is at least as large as the cost of the policy. Of course the benefits (and often some of the costs) usually occur in the future, and therefore must be expressed in present value terms. So, given a discount rate, it's all quite simple: calculate the present value of the benefits of a policy, subtract the present value of the costs, and see whether the difference (the Net Present Value, or NPV) is positive. And if one is comparing several alternative policies, choose the one with the highest NPV.¹

The student, however, may start to realize that the problem is in fact more complicated:

- First, we never really know what the benefits from reduced environmental damage will be, or even the amount that environmental damage will be reduced by a particular policy. Worse yet, we *can't know* with much precision what those benefits will be, even if we work very hard to find out. Take global warming. Modern meteorological science tells us that the relationships between greenhouse gas (GHG) concentrations, temperatures (regional or global) and climate patterns are inherently random. And even if we knew what those changes in temperatures and climate patterns are likely to be, we know even less about their economic and social impact, in part because we don't know how humans will adapt (e.g., by growing different crops or living in different areas). And if you think global warming is an unfair example because of the very long time horizons involved, take acid rain. Although virtually everyone would agree that the acidification of lakes and rivers – as well as the direct effects on human health – from unregulated nitrogen

oxide (NO_x) and sulfur oxide (SO_x) emissions is not a good thing, there is very little agreement as to just how bad it is.

- Second, we usually don't know what the current and future costs of a policy will be. In the case of a carbon tax, for example, we don't know how consumers and producers will respond, especially over the longer term. For example, to what extent will consumers use less fuel and buy more fuel-efficient cars and heating systems? And will producers develop and adopt more fuel-efficient technologies? Or in the case of NO_x and SO_x emissions, we know the current cost of scrubbers, but we don't know what their cost will be in the future, and how installing them will affect electricity prices and demand.
- And third, what discount rate (or rates) should be used to calculate present values? There is disagreement among economists over what is the "correct" rate that accounts properly for social time preferences and risk. And even if we settled on a conceptual notion of a "correct" rate, e.g., society's marginal rate of return on capital, there would still be considerable uncertainty over the actual numbers for current and future discount rates. (The marginal return on capital is difficult to measure, and its future evolution is inherently uncertain.) Furthermore, as we will see, discount rate uncertainty is itself a determinant of the "correct" effective rate that should be used for policy evaluation.

Of course the student might argue that there is nothing problematic about uncertainty over current and future benefits, costs and discount rates. If the student had taken a basic finance course, she would know that firms frequently make capital investment decisions in the face of similar uncertainties over the future cash flows from the investment, and must select discount rates subject to uncertainty over the correct risk premium. She might argue that firms typically base their investment decisions on the expected values of those cash flows, and environmental policy design can likewise be based on expected values.

The student might also argue that most public policy decisions must be made in the face of uncertainty. Possible changes in our Social Security or Medicare programs must be evaluated in the context of a broad set of uncertainties: over future changes in the demographic makeup of the country, changes in incomes, savings rates, and costs of living for different demographic groups, and changes in disease prevalence and medical costs, to name a few. Likewise, it is notoriously

¹ For a good textbook discussion of cost-benefit analysis applied to environmental policy, see Tietenberg (2006).

difficult to predict the effects of changes in tax policy on income, employment, and government tax revenues. Again, what's special about environmental policy?

I would counter that for many environmental problems, the uncertainties are greater and more crucial to policy design and evaluation. In particular, three important complications arise that are often crucial for environmental policy, but are usually much less important for most other private and public policy decisions.

1. The first complication is that environmental cost and benefit functions tend to be *highly nonlinear*. In other words, the damage likely to be caused by air or water pollution or by GHG emissions does not increase linearly with the level of pollution or emissions. Instead, the damage might be barely noticeable for low levels of pollution and then become severe or even catastrophic once some (uncertain) threshold is reached. Likewise, the cost of pollution abatement may be very low for low levels of abatement, but then become extremely high for higher or total abatement. This means that one cannot simply use expected values; the expected value of the cost or benefit function will be very different than the function of the expected value.

Furthermore, the precise shapes of the functions are unknown. This is particularly important if we believe that there is a threshold or “tipping point” at which the impact of a pollutant becomes extremely severe, but we don't know where that point is. For example, how large an increase in GHG concentrations – or in mean temperature – would it take for the consequences to be near-catastrophic? And at what point would over-fishing or habitat destruction lead to the collapse or extinction of a fish or animal population? A lack of answers to these questions suggests that environmental policy should be “precautionary,” in the sense of favoring earlier and more intense intervention. But just how “precautionary” should the policy be? Should countries agree, for example, to roll back their GHG emissions to 1930 levels? I do not mean to be facetious. Rather, I want to stress that uncertainty over the existence and/or position of a “tipping point” can be critical to policy timing and design.

2. The second complication is that environmental policies usually involve important *irreversibilities*, and those irreversibilities interact in a sometimes complicated way with uncertainty. There are two kinds of irreversibilities that are relevant for environmental policy, and they work in opposite directions.

First, policies aimed at reducing environmental degradation almost always impose sunk costs on society. These sunk costs can take the form of discrete investments (e.g., coal-burning utilities might be forced to install scrubbers), or they can take the form of expenditure flows (e.g., a price premium paid by a utility that has committed to burning low-sulfur coal). In either case, if future policy costs and benefits are uncertain, these sunk costs create an opportunity cost of adopting the policy, rather than waiting for more information about environmental impacts and their economic consequences. This implies that traditional cost-benefit analysis will be biased toward policy adoption.

Second, environmental damage is often partly or totally irreversible. For example, atmospheric accumulations of GHGs are long lasting; even if we were to drastically reduce GHG emissions, atmospheric concentration levels would take many years to fall. Likewise, the damage to ecosystems from higher global temperatures, acidified lakes and streams, or the clear-cutting of forests may be permanent. This means that adopting a policy now rather than waiting has a sunk benefit, i.e., a negative opportunity cost. This implies that traditional cost-benefit analysis will be biased against policy adoption.

How important are these irreversibilities and what are their implications for policy? The answers depend on the nature and extent of the uncertainties over costs and benefits, and how those uncertainties are likely to get resolved over time. The greater the current uncertainties, and the greater the rate at which they will be resolved, the greater will be the opportunity costs and benefits associated with policy adoption.

3. Third, unlike most capital investment projects and most other public policy problems, environmental policies often involve *very long time horizons*. While NPV calculations for firms' investments rarely go beyond 20 or 25 years, the costs and especially the benefits from an environmental policy can extend 100 years or more. The problems of global climate change and nuclear waste disposal are well-known examples with long time horizons, but there are others. For some forests and the ecosystems they contain, clear-cutting and other interventions can have consequences that extend for many decades. Likewise with chemical contaminations of land or water supplies. And the extinction of a species is, by definition, forever.

A long time horizon exacerbates the uncertainty over policy costs and benefits. It is hard enough to predict the impact of pollution, or the costs of abatement, five or ten years from now. Over a 50-year horizon, the uncertainties are much greater.

A long time horizon also makes discount rate uncertainty much more important. Suppose we are not sure whether the “correct” discount rate for evaluating a policy is 2 percent or 4 percent. (Many rates have been used in policy analysis; the OMB has suggested 3 percent and 7 percent for government regulatory analyses.) With a 2 percent discount rate, a \$100 benefit 50 years from now is worth about \$37 today, but with a 4 percent rate it is worth only \$14 today. If the \$100 benefit accrues 100 years from now, its present value is about \$14 with the 2 percent rate, but only \$2 with the 4 percent rate. Clearly with discount rates of 4 percent or more, it would be very hard to justify almost any policy that imposes costs today but yields benefits only 50 or 100 years in the future.

Uncertainty over future discount rates has an important implication for the choice of discount rate that we should use in practice – it makes that rate *lower* than any expected future discount rate. The expected present value of \$100 received T years from now is the expected value of $\$100/(1+R)^T$, where R is the (uncertain) discount rate (and $1/(1+R)^T$ is called the *discount factor*). But the expected value of $\$100/(1+R)^T$ is *greater than* $\$100/(1+R_e)^T$, where R_e is the expected value of R . Furthermore, the longer the time horizon (i.e., the larger is T), the greater is the difference between the expected value of $\$100/(1+R)^T$ and $\$100/(1+R_e)^T$. This means that the effective discount rate that should be used in a present value calculation is *less than* the expected (or average) discount rate.

The remainder of this paper discusses the sources and nature of the uncertainties that tend to arise in environmental economics, and how the policy implications of those uncertainties are shaped by the nonlinear nature of benefit and cost functions, by the irreversibilities that are often present, and by long time horizons. The next section explains why benefits and costs are inherently uncertain, and discusses the nonlinear characteristics of environmental benefit and cost functions. I focus in particular on uncertainty over possible “tipping points,” i.e., a threshold resulting in catastrophic environmental damage. Section 3 deals with irreversibilities and the opportunity costs and benefits they create, and Section 4 examines the implications of

long time horizons and discount rate uncertainty. The paper concludes by summarizing the lessons for policy design, and the areas where much more research is needed.

Despite its title, this paper is not intended to be a comprehensive survey of the many aspects of uncertainty in environmental economics, or the vast amount of recent and ongoing research on the topic. This is, after all, an article, not a book, and I have chosen to focus on those areas where I have a research interest, and where I believe some elucidation would be worthwhile.

2. The Uncertain Nature of Benefits and Costs

I claimed in the previous section that in comparison to many other public policy problems, environmental problems typically involve uncertainties that are greater and more crucial to policy design and evaluation. Why? As I explained above and will discuss in more detail below, environmental policy design must contend with highly nonlinear benefit and cost functions, irreversibilities, and long time horizons. But in addition, environmental problems usually involve three compounding levels of uncertainty – uncertainty over the underlying physical or ecological processes, uncertainty over the economic impacts of environmental change, and uncertainty over technological changes that might ameliorate those economic impacts and/or reduce the cost of limiting the environmental damage in the first place. (By “economic impacts,” I mean to include health impacts, lost consumer and producer surplus from degraded air, water, fisheries and other public goods, as well as lost output resulting directly from changes in climate, resource availability, etc.)

These compounding levels of uncertainty apply to both the benefit and cost sides of policy design and evaluation. For example, the future benefits of reducing GHG emissions depend first on the uncertain relationships between GHG emission levels, GHG concentrations, and resulting temperature distributions (as well as other climatic effects). But those benefits also depend on the uncertain economic impacts of changes in temperature distributions, as well as adaptation to climate changes (e.g., the use of new plant varieties or irrigation methods) that might reduce those economic consequences. The current and future costs of reducing anthropogenic climate change depend on the amount by which GHG emissions would need to be reduced, but that in turn depends on the uncertain relationships between GHG emission levels, concentrations, and climate effects, and the uncertain economic impacts of the climate effects. (It is the economic impacts, broadly construed, that is the policy concern in the first place.) And those costs also

depend on unpredictable technological advances in energy conservation, cheaper non-polluting energy sources, and methods of carbon sequestration.

These uncertainties and their effects on policy are magnified by the often nonlinear nature of benefit and cost functions. Uncertainties over benefits and costs are manifested not only in the form of parameter uncertainty (e.g., uncertainty over the elasticity of emissions with respect to a tax rate on emissions), but also in the form of uncertainty over the shapes of the (nonlinear) benefit and cost functions (e.g., uncertainty over how that elasticity falls as the tax rate is increased). As discussed below, the problem becomes especially severe when there are “tipping points” (i.e., when at some level of environmental damage the consequences become near-catastrophic), but we don’t know what that point is.

To make the discussion over benefits and costs more concrete, I will focus on global warming, for which the policy implications of uncertainty have been studied extensively (but certainly not exhaustively). Later I will turn to the nature and implications of uncertainty for a very different environmental policy problem – the regulation of a fishery.

2.1. Uncertainty over Benefits

The point of environmental policy is to bring human exploitation of environmental assets closer to socially optimal levels, thereby creating social benefits. We might impose a carbon tax to reduce the future economic impact of global warming because we expect the benefits (a reduced economic impact from reduced warming) to outweigh the costs of the policy. But how soon should we impose a carbon tax, and how large should it be? Putting aside costs for the moment, the answers depend on the specific benefit functions, i.e., on how the benefits from the tax vary with its size. And the answers also depend on nature and extent of uncertainty over those benefits.

The benefits over the next 100 years from reducing GHG emissions depend on (1) expected GHG emission levels absent abatement; (2) how rapidly atmospheric GHG concentrations will grow given these emission levels; (3) how higher GHG concentrations will affect global temperatures; and (4) how large an economic impact we should expect from higher temperatures.² The uncertainties are substantial for each step in this chain. Combining the first

² I am simplifying the problem for expositional purposes. For example, GHGs include methane as well as CO₂, and the mix matters. Likewise, it is not simply the change in mean temperature that matters, but also the geographic

three steps, if we take the Intergovernmental Panel on Climate Change (IPCC) report (2001) as a benchmark, the projected increase in mean temperature by the end of this century absent an abatement policy ranges from 1.4°C to 5.8°C. This large range should not be surprising. GHG emission levels are hard to predict because they depend on uncertain economic growth and energy intensities, and energy intensities in turn depend on unpredictable changes in energy prices and technologies.³ And even if emission levels were known, changes in GHG concentrations and changes in temperature depend on a complex physical system (with feedback loops that might be positive or negative) that is poorly understood.

The economic impact of global warming is even harder to predict. A temperature increase in the middle of the IPCC range – say, 2.5°C – would likely be beneficial for some regions of the world (e.g., Russia and Canada), harmful for other regions (e.g., India, Bangladesh, and Southeast Asia), and might have little or no aggregate net impact on other regions (e.g., the U.S.). Given the long time horizon involved, we would expect people to adapt to changing temperatures and changing climates generally (e.g., by planting and developing different crops, and by moving from warm and low-lying areas to cooler and higher ones). Studies have shown that potential adaptation is clearly important (e.g., Mendelsohn, Nordhaus, and Shaw (1994)), but it is still difficult to predict to what extent people will adapt, and at what cost.

A number of studies have attempted to assess the sources and extent of uncertainty over the benefits from reducing GHG emissions. One of the earliest was Nordhaus (1994b), who obtained estimates of the percentage loss in gross world product from a survey of natural scientists and economists. Roughgarden and Schneider (1999) then used the Nordhaus survey results, along with other survey evidence, to back out confidence intervals for a damage function. Pizer (1999) developed an integrated climate-economy model, and used it to both assess parameter uncertainty and study its policy implications. His model has nineteen parameters. Six parameters describe economic activity, and are estimated econometrically using historical data,

distribution of temperature changes. And of course economic impacts can be complex and vary substantially across geographical regions.

³ Kelly and Kolstad (2001), for example, point out that future GHG emissions depend critically on population and productivity growth, which are uncertain. Lower growth would make global warming less of a problem. On the other hand, lower productivity growth implies a lower discount rate, which, as Pizer (1999) shows, increases the present value of the future benefits from abatement.

yielding a joint distribution. Uncertainty over the remaining thirteen parameters (describing emissions growth, GHG retention and decay rates, control costs, and long-term population and productivity growth) is based on “subjective analysis.” Other studies likewise assess uncertainty using subjective analysis and expert opinion; for an overview, see Heal and Kriström (2002).

One could argue that what really matters is the possibility and consequences of temperature increases at or above the upper end of the IPCC range. How likely is a 6°C temperature increase, and what impact would it have? One possible scenario (absent the probability estimates) was depicted by Gore (2006) – large parts of New York City would be under water. (Could New Yorkers adapt? Perhaps.) Clearly, the issue is whether for some plausible increase in temperature, things could get really bad, to the point where adaptation couldn’t compensate. In other words, is there a tipping point in the benefit function, and if so, where is it? Unfortunately, we don’t know.

As of now, very little work has been done to assess the probabilities of catastrophic climate change, or to estimate the point (e.g., the change in mean temperature) at which a catastrophic outcome becomes likely. People have widely differing (and often strongly held) opinions about the likelihood and makings of a catastrophic outcome, but this is just symptomatic of how little we know. There have been studies of the policy implications of catastrophic outcomes, which I will discuss later, but those studies take the outcome and its characteristics as given.⁴

This problem, which I will call “tipping point uncertainty,” is not unique to global warming. For example, studies of toxic waste disposal suggest that points could be reached that are catastrophic for land and water use in localized areas. But again, there is considerable uncertainty as to where those points are.

In summary, what do we know about uncertainty over the benefits of environmental policy? As I have tried to show, in the case of global warming, which has been studied quite extensively, we know very little. That is, we know that there is a good deal of uncertainty, but we are hardly

⁴ For example, Pizer (2003) modifies the DICE model developed by Nordhaus (1994a), replacing the original quadratic relationship between economic damage and temperature change with a function that is much more convex (and varies the degree of convexity). But his focus is on whether price versus quantity is the more effective policy instrument. As he says, “Unfortunately, there is little empirical information concerning either the degree of steepness or the point where the steepness begins.”

able to quantify it, especially when it comes to tipping points. We would come to a similar conclusion for environmental problems that have received less attention (e.g., acid rain and toxic waste). Given that uncertainty clearly matters (as I will explain in more detail below), we seem to have our work cut out for us.

2.2. Uncertainty over Costs

For some environmental problems, particularly those with more limited time horizons, policy costs are better understood and subject to less uncertainty than are the benefits. For example, we have years of experience with limits on SO_x and NO_x emissions from coal-burning power plants. We know the cost (and effectiveness) of scrubbers and of substituting low-sulfur coal, and we can also infer costs from prices of tradable emission allowances. And while the cost function for emission controls is nonlinear (the cost of an incremental reduction in emissions rises rapidly as the emission level becomes very low), there is no “tipping point” problem.

For other environmental problems, however, cost uncertainty can be quite severe. Once again, global warming is a good example. A carbon tax may be the preferred instrument for reducing carbon dioxide (CO₂) emissions, but how large a tax would it take to reduce those emissions in the U.S. by, say, 20 percent?⁵ The answer depends on how responsive fossil fuel demand would be to tax-induced price changes, and this will vary from sector to sector (e.g., transportation versus residential heating). The responsiveness of fossil fuel demand in any sector to price changes in turn depends on the long-run price elasticity of energy demand in the sector, and the long-run elasticity of substitution between fossil and non-fossil energy sources. We have a reasonable understanding of energy demand elasticities, but our knowledge of the long-run elasticity of substitution between fossil and non-fossil energy sources is tenuous at best. The reason is that the ability to substitute from, say, coal to wind power for electricity production depends on the cost and availability of the latter (and that cost will depend in part on its environmental impact). Today, that substitutability is quite limited; what it would be 20 or 50 years from now depends in large part on technological change, which is, again, hard to predict.

⁵ We would also need to know how large a tax it would take to reduce emissions in other countries, and for developing countries like China and India, the uncertainties are far greater than for the OECD. Then there is the cost of free riding. How large a tax would be needed for the OECD countries if the larger developing countries did not also agree to reduce emissions?

Another way to look at this is to ask whether (and how fast) the cost difference between fossil-based and alternative energy supplies will converge. Solar, wind, and biomass costs have fallen over the past twenty years, but the differences are still very large. There is reason to think that over the next fifty years, those differences may decline or even disappear as fossil fuel prices rise (because of depletion) and costs of alternative energy sources fall (because of technological change and economies of scale). Chakravorty, Roumasset, and Tse (1997) have developed an empirical model of this process that takes into account the depletion of the potentially discoverable reserves of fossil fuels. They project that the world will move towards wide-scale use of solar energy, so that “90 percent of the world’s coal will never be used,” and global temperature will rise only about 1.5 to 2°C over the next century, making the cost of mitigation low indeed. However, since standard errors are missing, one cannot attach confidence intervals to these optimistic forecasts. Other forecasts suggest starkly different results, with fossil fuels remaining the predominant energy source in the absence of major policy changes.

The basic problem is that with long time horizons, policy costs depend on technological change, which is inherently difficult to predict. In fact, it is difficult to even characterize the uncertainty over technological change. Thus, for global warming and other long-horizon environmental problems, the cost side of policy is subject to uncertainty over the uncertainty.

2.3. Implications for Policy Design.

Uncertainties over benefits and costs can affect policy design in at least three fundamental ways. First, they can affect the optimal choice of *policy instrument*, i.e., whether pollution is best controlled through a price-based instrument (e.g., an emissions tax) or a quantity-based instrument (e.g., an emissions quota). Second, they can affect the optimal *policy intensity*, e.g., the optimal size of the tax or the optimal level of abatement. Third, they can affect the optimal *timing of policy implementation*, i.e., whether it is best to put an emissions tax in place now, or wait several years (and thereby reduce some of the uncertainty).

Choice of Policy Instrument. The implications of uncertainty for the optimal choice of policy instrument has been studied extensively, beginning with the seminal paper by Weitzman (1974), who showed that in the presence of cost uncertainty, whether a price-based instrument or a quantity-based instrument is best depends on the relative slopes of the marginal benefit function and marginal cost function. If the marginal benefit function is steeply sloped but the

marginal cost function is relatively flat, a quantity-based instrument (e.g., an emissions quota) is preferable: an error in the amount of emissions can be quite costly, but not so for an error in the cost of the emissions reduction. The opposite would be the case if the marginal cost function is steeply sloped and the marginal benefit function is flat. Of course in a world of certainty, either instrument will be equally effective. If there is substantial uncertainty and the slopes of the marginal benefit and cost functions differ considerably, the choice of instrument can be crucial.

Weitman's original result has been extended in a number of directions. For example, Stavins (1996) showed that a positive correlation between marginal benefits and marginal costs pushes the optimal choice towards a quantity instrument. Often, however, there is no need to choose exclusively between a price and a quantity instrument. A number of studies have shown that in the presence of uncertainty, "hybrid" policies that combine both instruments generally dominate the use of a single instrument (e.g., Roberts and Spence (1976), Weitzman (1978), and in the context of climate change policy, Pizer (2002) and Jacoby and Ellerman (2004)). The optimal design of a hybrid policy depends not only on the shapes of the benefit and cost functions, but also on the nature and extent of the uncertainties. Our lack of knowledge of the uncertainties (and often the shapes of the functions) means that policy design will be sub-optimal at best; much more work is needed to get a better understanding of just how sub-optimal.

Policy Intensity. Uncertainties over benefits and costs can also affect the optimal policy intensity, i.e., the size of an emissions tax or the amount of abatement that should be mandated. If there are no irreversibilities (which are discussed below), then in many cases uncertainty will lead to a *lower* policy intensity.

Suppose the policy instrument is an emissions quota, which we can think of as an amount of abatement between 0 and 100 percent. The greater the abatement, the greater will be the resulting benefit, and the greater will be the cost. However, in many cases we would expect the benefit function to be concave: increasing the abatement from 0 to 10 percent will have a large incremental benefit, whereas increasing the abatement from 90 to 100 percent is likely to have a much smaller incremental benefit. (This would be the case, for example, if a high level of pollution has serious health effects, so that there is a substantial benefit even from reducing emissions by 10 percent, whereas a low level of pollution has negligible health effects, so that there is very little added benefit from going from a 90 percent to 100 percent emissions reduction.) Likewise, the cost of abatement is usually convex: going from no abatement to a 10

percent reduction in emissions is likely to be much less costly than going from a 90 percent reduction to a 100 percent reduction. If the uncertainties over benefits and costs are proportional to the amount of abatement, it follows that the optimal amount of abatement (which equates expected benefits with expected costs) will be lower than in the absence of uncertainty.⁶ On the other hand, if there is no cost uncertainty, and the uncertainty over benefits is proportional to the level of emissions (as opposed to the abatement level), the optimal abatement would be higher than in the absence of uncertainty.

One might ask whether the possibility that little or no abatement would lead to a catastrophic outcome might increase the optimal abatement level. It is important not to confuse the *expected* benefit from abatement with the *variance* of that benefit. The possibility of a catastrophe will likely increase the expected benefit from any amount of abatement, which, other things equal, implies a greater optimal abatement level. But if we are interested in the effect of uncertainty alone, we must keep the expected benefit fixed as we increase the variance of the benefit.

Policy Timing. Uncertainty can also affect the optimal timing of policy implementation – but only if there are sunk costs of implementing the policy and/or the environmental damage from having no policy in place is at least partly irreversible. Depending on the particular situation, it may be optimal to defer the implementation of a policy until we learn more about benefits and costs, or to speed up the implementation to avoid irreversible damage. I address this implication of uncertainty below as I discuss irreversibilities.

3. Irreversibilities

It has been understood for many years that environmental damage can be irreversible, and that this can lead to a more “conservationist” policy than would be optimal otherwise. To my knowledge, the earliest economic studies to make this point are Arrow and Fisher (1974) and Henry (1974). But thanks to Joni Mitchell, even non-economists know that if we “pave paradise and put up a parking lot,” paradise may be gone forever. And if the value of paradise to future generations is uncertain, the benefit from protecting it today should include an “option value,”

⁶ Suppose the cost of abatement is $C(A) = c[(1+\varepsilon)A]^2$, where c is a constant, A is the percentage abatement, and ε is random and equal to either -1 or $+1$, each with probability $\frac{1}{2}$, so the expected value of ε is 0 . Using this expected value, the cost of abatement is cA^2 . However, the expected cost of abatement is $\frac{1}{2}c(0 + 4A^2) = 2cA^2$. The uncertainty over ε increases the expected cost of abatement, making the optimal amount of abatement smaller.

which pushes the cost-benefit calculation towards protection.⁷ Thus it might make sense to restrict commercial development in a wilderness area, even if *today* very few people care at all about that wilderness or have any desire to visit it. After all, if it becomes clear over the next few decades that very few people will *ever* care much about that wilderness, we will still have the option to pave it over. But if we exercise the option and pave over the wilderness today, we will lose the flexibility that the option gave us. If the social value of the wilderness rises sharply over time, the social loss from having exercised our option prematurely could be great.

There is a second kind of irreversibility, however, that works in the opposite direction: Protecting paradise can impose sunk costs on society. If paradise is the wilderness area, restricting development and human access would imply a (permanently) foregone flow of wage and consumption benefits (e.g., from a large ski resort). If paradise is clean air and water, protecting it could imply discrete sunk cost investments in abatement equipment, and/or an ongoing flow of sunk costs for more expensive production processes. In both cases, these costs are permanent – we cannot recapture them in the future should we decide that clean air and water are less important than we had originally thought. The point is that this kind of irreversibility would lead to policies that are *less* “conservationist” than they would be otherwise.

The implications of irreversibilities have been explored in a large and growing literature.⁸ The important questions relate to the conditions under which irreversibilities matter for policy, and the extent to which they matter.

3.1. When Do Irreversibilities Matter?

It is important to stress that *these irreversibilities only matter if there is uncertainty*. To see why, suppose that today we know *precisely* how society will value a pristine wilderness area every year over the next 200 years, and we also know what would be the annual flow of profits,

⁷ “Option value” and “quasi-option value” have been used in the environmental economics literature in different and sometimes confusing ways. Arrow and Fisher (1974) referred to the value of waiting when environmental damage is irreversible as “quasi-option value.” The term “quasi-option value” has been adopted by other authors, who use “option value” to refer to the value of delay because of risk aversion. See, e.g., Conrad (1980), Freeman (1984), and Hanemann (1989). I will use the term “option value” to refer to any opportunity costs or benefits resulting from irreversibilities and uncertainty. This is consistent with the real options literature, e.g., Dixit and Pindyck (1994), and with modern textbook usage, e.g., Tietenberg (2006).

⁸ Fisher and Hanemann (1990) and Gollier et al (2000) extend the results of Arrow and Fisher (1974) and Henry (1974), as do Scheinkman and Zariphopoulou (2001), but using a continuous-time model. Kolstad (1996b) and Ulph and Ulph (1997) address conditions under which irreversibilities matter. Implications for the timing of policy adoption are studied in Pindyck (2000, 2002). These are examples; for a survey, see Heal and Kriström (2002).

wages, and consumer surplus over that same period from the conversion of that wilderness into a large commercial resort. Finally, suppose we know the correct discount rates that would apply over that 200 year period. We could then calculate present values and do a simple cost-benefit comparison of the pristine wilderness and the commercial resort. If the comparison favored the resort, we could allow commercial development, knowing that nothing would change in the future that would lead to regret and a desire to “un-do” the loss of the wilderness. Likewise, if the comparison favored the wilderness, we could prevent commercial development, knowing that nothing would change that would lead to regret over the loss of surplus from a resort. Irreversibility would be irrelevant.

The environmental economics literature has also been concerned with the question of when irreversibility matters in terms of affecting current decisions, even if there is uncertainty (e.g., Kolstad (1996b) and Ulph and Ulph (1997)). The short answer is that irreversibility will affect current decisions if it would constrain future behavior under plausible outcomes. Consider global warming and suppose that benefits and costs are completely linear and there is no risk aversion. Then the relevant constraint is that while we could (at great cost) have zero emissions in the future, we cannot have negative emissions.⁹ If there is a non-negligible chance that we might want to have negative emissions in the future, the irreversibility constraint might bind, and should lead to lower emissions today. However, irreversibility could matter even if there is no chance we would want negative emissions. Suppose, for example, that the cost of reducing emissions rises more than proportionally with the amount of the reduction. In that case, we would lower emissions today by an amount of, say, x , so that we could avoid having to lower them in the future by $2x$ at a cost that is more than twice as large.

The presence of uncertainty can affect policy even if there are no irreversibilities, but in a more limited way than if irreversibilities are present. As explained earlier, uncertainty can affect policy because of nonlinear cost or benefit functions, or discount rate uncertainty. For example, suppose the environmental impact of air pollution was completely reversible so that each year society could impose a quota on the pollutant concentration, at a cost proportional to the quota. Suppose also that each year the *actual* concentration equaled the quota plus a zero-mean proportional error. Finally, suppose (realistically) that the adverse health effect of pollution rises

⁹ This is not precisely correct. With near-zero emissions and a carbon sequestration program, effective emissions could be made negative. But as a practical matter, negative emissions are hard to imagine.

more than proportionally with the concentration, so that the benefit loss from a 5-percent unexpected increase in the concentration above the mandated quota is greater than the benefit gain from a 5-percent unexpected decrease. Compared to the case of no random errors, the optimal quota would then be lower.

But now suppose the atmospheric accumulation of this pollutant is very long-lived, so that emissions could be reduced, but not the accumulation of past emissions. Because each year's actual emissions is partly random, future concentrations are likewise random, and the variance of those concentrations will grow with the time horizon. In this case, the policy implications of uncertainty are much greater. The reason is that the range of possible future concentrations is far greater, and should those concentrations turn out to be very large, we will be stuck with them – we cannot “un-do” the emissions quotas of the past which *ex post* have turned out to be too high. As a result, the optimal policy would call for a smaller quota than would be the case if accumulations of the pollutant were very short-lived.

It is important to be clear about why irreversibility is key. Returning to our example, the fact that accumulations of the pollutant are long-lived creates a possibility of severe regret which is not offset by the possibility of very little pollution. A “bad-news principle” is at work here: If future concentrations of the pollutant turn out to be less than expected (“good news”), or if the economic and health impact of the pollution turns out to be less than expected (more “good news”) we could always relax the quota. But if future concentrations and/or the economic and health impact turn out to be greater than expected (“bad news”), there is little we can do to correct the situation. Even if we make our quota much more stringent, it will take many years for the concentrations to fall. It is this possibility of “bad news” that affects current policy, and the greater the uncertainty, the greater that possibility.¹⁰

In the example above, I focused on the irreversibility of environmental damage, but the optimal policy would also be affected by the irreversibility inherent in the sunk costs of abatement. To see this, suppose that atmospheric accumulations of this pollutant were very short-lived, but to abate, companies would have to install long-lived equipment. Then, if we later learn that the economic and health impact of the pollutant are much less than expected, we will be stuck with the equipment – we cannot “un-do” the emission quotas of the past which *ex*

¹⁰ This point was introduced by Bernanke (1983), and is discussed in detail in Dixit and Pindyck (1994). The extension to environmental policy (and the “good news principle” discussed below) is in Pindyck (2002).

post have turned out to be too small. In effect, a “good news principle” would apply: It is the possibility that the harm from the pollutant will turn out to be less than expected (“good news”) that would lead to regret. As a result, the optimal policy would call for *lower* mandated abatement than would be the case if the capital equipment used to abate were very short-lived.

3.2. How Do Irreversibilities Affect Policy?

We have seen that irreversibilities can interact with uncertainty to affect current policy, sometimes making it more and sometimes less “conservationist.” But how important are irreversibilities, and what is their overall effect on policy? If we ignore them, would we be led seriously astray when designing policies to deal with, say, global warming or toxic waste disposal? And would taking them into account make us more or less “conservationist?”

As one would expect, the answers depend on the particular policy problem. Unfortunately, there are very few policy problems for which the effects of irreversibilities have been studied in any detail. Not surprisingly, the problem that has received the most attention in this context is global warming, a problem for which we have a good sense of the irreversibilities involved, and as discussed earlier, the uncertainties are considerable. But even here, we have no clear answers as to the importance of irreversibilities and the overall effect they have on policy. The reason is that to determine the optimal abatement policy for any realistic climate-economy model, one must solve (analytically or numerically) a complex stochastic dynamic programming problem. Thus, researchers have had to make strong simplifying assumptions, and different sets of simplifying assumptions have led to quite different results. This is best seen by summarizing the approaches and results of several of the studies published over the past decade.

One of the most common simplifying assumptions is to limit time to two periods, so the question becomes how much emissions should be reduced now versus in the future. This was the approach used by Kolstad (1996b) in one of the earliest studies of the opposing irreversibilities – long-lasting GHG concentrations versus long-lived abatement capital – involved with global warming. He examined how the prospect of more information about the economic impact of GHGs in the second period would affect sunk expenditures on abatement in the first period. Not surprisingly, both irreversibilities can matter, but the net effect on first-period policy depends on the relative decay or depreciation rates of GHG concentrations and abatement capital, and on the expected benefit of abatement.

In another study, Kolstad (1996a) adapts the Nordhaus (1994a) DICE model by introducing uncertainty – which is reduced as learning occurs – over whether global warming will be a “big” or “little” problem. He calculates the optimal abatement policy over 20 ten-year periods, and finds that only the irreversibility associated with abatement capital matters, so that near-term policy should be less conservationist. The reason is that temperature change from GHG build-up is sufficiently slow so that emissions could always be slightly reduced in the future. Using a two-period model, also with functional forms and parameters taken from the DICE model, Fisher and Narain (2003) likewise find that the investment irreversibility effect is much larger than the GHG irreversibility effect, so that uncertainty over the impact of climate change leads to a reduction in first-period abatement.

In Pindyck (1998, 2000, 2002), I develop continuous-time models in which there is ongoing uncertainty over both the benefits from reduced GHG concentrations and the evolution of those concentrations, concentrations are long-lived, and there are sunk costs of policy adoption. I focus on the timing and size of a single permanent reduction in emissions (with the sunk cost of abatement a quadratic function of the size of the reduction). I find that for a “reasonable” range of parameters, either kind of uncertainty leads to a higher threshold for policy adoption and a less stringent reduction in emissions. The reason is that policy adoption commits to a reduction in the entire trajectory of future emissions at a large sunk cost, while inaction over any short time interval only involves continued emissions over that interval.¹¹

The studies cited above, however, do not consider possible catastrophic impacts of GHG accumulations. Those that do consider catastrophic impacts find that they can lead to earlier and more stringent abatement policies, but only if the likelihood of a catastrophe is strongly linked to the GHG concentration. This ambiguous result is shown clearly in Clarke and Reed (1994), who develop a simple but revealing theoretical model in which consumption causes the emission of an accumulating pollutant (so that reducing emissions requires consuming less), and the catastrophe is a random (Poisson) arrival in which welfare is reduced permanently to zero. The

¹¹ Newell and Pizer (2003b) and Pizer (2005) also develop dynamic models to study optimal policies when GHG concentrations are long-lived (but abatement capital is not). They focus on the choice of policy instrument, and show that despite the irreversibility, price-based policies (e.g., taxes) are strongly preferred to quantity-based policies (e.g., tradable emission allowances). Kelly and Kolstad (1999) examine the implications of learning (e.g., about parameters of the benefit function). They show that the ability to learn can imply larger emissions now, because a greater GHG concentration provides more information about parameter values. Gollier, Jullien, and Treich (2000) use a two-period model with irreversibility and learning to show how effects of uncertainty depend on the shape of the representative consumer’s utility function.

hazard rate could be a constant, or it could be an increasing function of the stock of pollutant. They show that if the hazard rate is a constant, we should accept *more* pollution now, because we will all be dead (or at least have zero welfare) at some point in the future that is unrelated to how much we pollute. On the other hand, if the hazard rate is sufficiently increasing with the stock of pollutant, we should pollute *less* now, because that will lengthen our expected remaining time on the planet. If the hazard rate increases only slowly with the stock of pollutant, the impact on the current level of pollution could go either way.

In related work, Tsur and Zemel (1996) derive optimal abatement policies using a model in which an undesirable (possibly catastrophic) event is triggered when the stock of pollutant reaches a critical level – but that critical level is unknown. This kind of event uncertainty leads to strongly cautionary behavior: The optimal emissions policy keeps the stock of pollutant within a fixed interval. Finally, Pizer (2003) studies catastrophic outcomes by replacing the quadratic relationship between economic damage and temperature change in the Nordhaus (1994a) DICE model with a more convex function, and varies the degree of convexity. But there are no sunk costs of adoption, and his focus is on whether price versus quantity is the more effective policy instrument.

So where does this leave us in terms of the policy implications for global warming? The irreversibilities are clear, and the uncertainties are great. Should we adopt a stringent emissions reduction policy now, despite its cost, or go slowly and wait to learn more about the rate of global warming and its likely economic impact? To my knowledge, research to date does not give us the answer. Those studies cited above that ignore possible catastrophic impacts provide some evidence that we should move slowly. Those studies that do consider the possibility of catastrophic impacts suggest a more stringent emissions policy, but the catastrophic impacts in those studies are more or less assumed, rather than inferred from empirical evidence. Once again, we have a good understanding of the economic theory, but a poor understanding of its implementation in practice.

3.3. Renewable Resource Management

So far, I have characterized environmental irreversibility in terms of a stock externality – emissions cause the build-up of a pollutant (e.g., GHGs), and it takes a long time for the stock of pollutant to dissipate. But environmental irreversibilities can arise in other ways, such as species

extinction or permanent loss of a wilderness area. An interesting area of environmental economics in which irreversibilities and uncertainty arise and interact in ways quite different from, say, global warming, is the management of a renewable resource such as a fishery.

The regulation of a fishery can be a complex problem even if the resource has a single owner. The problem becomes more complex (and of greater interest to environmental economists) when the resource is common property (i.e., there is open access). In both cases, the problem is how to regulate extraction (fishing) so as to maximize the economic value of the resource. There is a large literature on this topic, but most of the early work assumed that the growth function for the resource stock is known and deterministic, so that the dynamics of the stock can be described by a simple differential equation in which the rate of change of the stock equals the growth function minus the extraction rate. Thus, left to its own, the stock of fish will increase to some natural carrying capacity, but over-fishing will lead to a stock below the optimal level, or even drive the stock permanently to zero.¹² The environmental irreversibility arises because it may take some time for the stock to recover from over-fishing, and it will never recover if it is driven to zero.

Uncertainty enters into the problem because the resource growth function is in fact stochastic, and sometimes highly stochastic. In other words, the rate of change of the resource stock is not a deterministic function of the current stock level minus the rate of extraction. Instead, the stock dynamics must be described by a stochastic differential equation. Furthermore, we often can't observe the actual resource stock, but can only estimate its value subject to error. The optimal resource management problem then becomes a problem in stochastic dynamic programming, and its solution can be quite complicated.

How do stochastic fluctuations in resource growth affect the optimal regulated extraction rate? As with the global warming problem discussed above, there is no clear answer. Depending on the particular growth function and the characteristics of the extraction cost and resource demand functions, stochastic fluctuations could lead to a higher or lower regulated extraction rate. But there is an analogy with global warming. As a catastrophic outcome becomes a more distinct possibility, the optimal policy becomes more conservationist, i.e., an increase in the volatility of stochastic fluctuations in the resource stock leads to a reduction in the extraction

¹² For an excellent textbook treatment of renewable resource management and review of the early literature, see Clark (1990).

rate. The reason is that as the stock becomes smaller, there is a greater chance for a stochastic decline in the stock to (irreversibly) drive the resource to extinction.¹³

4. Discounting Over Long Time Horizons

The role of uncertainty in policy design is especially important for environmental problems that involve long time horizons. First, it is difficult enough to predict the costs and benefits over the next decade of reducing air or water pollution. Over a 50-year horizon, unpredictable technological change, changes in land and water use, and population shifts make the uncertainties over policy costs and benefits far greater. Second, discount rates are inherently uncertain, and a long time horizon makes that uncertainty extremely important. If the discount rate is “high” it will be difficult to justify almost any policy that imposes costs on society today but yields benefits only 50 to 100 years from now, so the size of the discount rate can be the make or break factor in policy evaluation. Third, uncertainty over future discount rates impacts the choice of discount rate that we should actually use in practice – it makes that rate *lower* than any expected future rate. I have already discussed the first and second of these three aspects of long-horizon uncertainty, so I will focus here on the third aspect – the fact that discount rate uncertainty reduces the effective discount rate that should be used for policy evaluation.¹⁴

As explained earlier, this is a straightforward implication of the fact that the discount factor is $1/(1+R)^T$, so with R uncertain, the *expected* discount factor is greater than the discount factor calculated using the expected value of R . A simple example might help to clarify this. Suppose we want to evaluate the expected present value of a \$100 benefit to be received 100 years from now, but we believe that the “correct” discount rate over the entire 100 years will turn out to be either 0 or 10 percent, each with probability $\frac{1}{2}$. If we apply the expected value of the discount rate, i.e., 5 percent, the \$100 future benefit will have a present value of less than \$1. But in fact the expected present value of the benefit is much higher than \$1. If the true discount rate turns out to be 0, the present value would be \$100, but if the discount rate turns out to be 10 percent,

¹³ I make no attempt to survey the literature on renewable resource management under uncertainty, but examples include Pindyck (1984) and, more recently, Singh, Weninger, and Doyle (2006). Weitzman (2002) shows that even with severe stochastic fluctuations in the resource stock, landing fees dominate quotas as a policy instrument.

¹⁴ By “effective discount rate,” I mean the single rate that could be used in place of the range of possible future rates. The fact that the effective discount rate is reduced by uncertainty over future rates has been discussed in detail by Weitzman (1998) and Newell and Pizer (2003). To my knowledge, this result was first formally demonstrated by Dybvig, Ingersoll, and Ross (1996), although in a different context.

the present value would be close to zero. Thus the expected present value of the \$100 benefit is $\frac{1}{2}(\$100) + \frac{1}{2}(\$0) = \$50$. Now ask what single discount rate when applied to a \$100 benefit received 100 years from now would yield a present value of \$50? The answer is about 0.7 percent. Thus even though the expected value of the discount rate is 5 percent, the uncertainty (it could in fact be either 0 or 10 percent) implies an *effective* discount rate of less than 1 percent.

Thus discount rate uncertainty reduces the effective discount rate that should be used to calculate present values. Of course the size of the reduction depends on the extent of the uncertainty. In the example I just gave, if the expected value of the discount rate were 5 percent but the range for the actual rate were between 4 percent and 6 percent, the effective rate would be only slightly less than 5 percent (4.58 percent, to be precise). This just means that understanding the nature and extent of discount rate uncertainty is a crucial step in doing the actual discounting, and thereby evaluating a policy that is expected to yield long-term benefits.

One approach to evaluating discount rate uncertainty is to estimate the parameters of the stochastic process followed by an appropriate interest rate. To do this, one needs a very long historical time series for the interest rate, which is problematic. Newell and Pizer (2003a) used this “reduced-form” approach by estimating an autoregressive equation for the interest rate using some 200 years of data on government bond rates. Putting the quality of the data aside, they found uncertainty to have a substantial effect on the discount rate for horizons of 100 years or more. For example, depending on the particular estimated interest rate equation, they find that the value today of \$100 to be received 100 (200) years from now is only \$1.98 (\$0.04) using a constant 4 percent discount rate, but is between \$2.61 and \$5.09 (\$0.10 and \$1.54) when uncertainty is taken into account.

Another approach is to note that what really matters are the tails of the distribution for future interest rates, and to try to estimate that part of the distribution. Over the long run, changes in real interest rates will come from changes in real economic growth, so what we want is the distribution for future real growth rates, and in particular, the tails of that distribution. For example, the possibility of rare disasters (not necessarily environmental in origin) would yield “fat tails” for the distribution of future real interest rates, and thereby considerably reduce the current effective discount rate. Barro (2005) extrapolates from the wars and depressions experienced by a variety of countries over the past century, and shows that possible disasters can

make the effective discount rate close to zero.¹⁵ However, the number of wars and depressions is limited, making it difficult to pinpoint the effective discount rate.¹⁶

What discount rate, then, should be used should be used to evaluate costs and benefits over long time horizons? The studies mentioned above do not give us a clear answer. However, they do show that the correct rate should decline over the horizon and for the distant future is probably well below 2 percent, which is lower than the rates often used for environmental cost-benefit analysis. Thus costly environmental policies whose benefits will come 100 years from now may indeed be justifiable.¹⁷

5. Conclusions.

Uncertainty is central to environmental policy. For most environmental problems, we have very limited knowledge of the underlying physical or ecological processes, the economic impacts of environmental change, and the possible technological changes that might occur and that would ameliorate the economic impacts and/or reduce policy costs. If costs and benefits were linear and both environmental damage and policy costs were reversible, these uncertainties would not complicate matter much; policies could be based on expected values of costs and benefits at each point in the future. But as I have tried to show, cost and benefit functions tend to be highly nonlinear, and both environmental damage and policy costs are often irreversible. As a result, it can be misleading to base policy on expected values of costs and benefits.

Then how should we base policy? Unfortunately, there are no simple guidelines or rules of thumb for adapting environmental policies to uncertainty, at least that I am aware of. As we have seen, the irreversibilities associated with environmental damage and policy costs work in

¹⁵ Barro's objective is to explain well-known asset-pricing puzzles, of which a near-zero real risk-free interest rate is one. (Another is the large observed risk premium on the market, which seems only to be consistent with an unreasonably high index of risk aversion on the part of investors.) Weitzman (2006) takes another approach to the same set of puzzles. He shows that parameter uncertainty (in terms of both the mean and the variance) with respect to the structure of the economy will lead to fat tails, which likewise imply near-zero real risk-free interest rates, and explain the large observed market risk premium.

¹⁶ In an engaging book, Posner (2004) examines a wide range of possible catastrophic events ranging from flu pandemics to nuclear war to an asteroid hitting the earth. Although his estimates of the probabilities are largely subjective, the number of possible events is large enough and the economic impacts severe enough to make one think that Barro's study is overly optimistic, and the left tail of the real growth rate distribution is quite fat indeed.

¹⁷ Also, there is evidence that consumers' subjective discount rates are much higher for short-run than long-run outcomes. For an overview of behavioral approaches to discount rates, see Frederick et. al. (2002). Weitzman (2001) also provides evidence of higher short-run than long-run rates, and shows that economists differ widely in their opinions regarding discount rates.

opposite directions, so their net effect tends to be model-specific. Recent studies of global warming that ignore catastrophic impacts show a net effect that favors waiting over early action. That result can be turned around in models in which the health and economic impact rises at a sharply increasing rate with the amount of pollutant, e.g., if at some level the pollutant could have a catastrophic impact. This would especially be the case if the point at which a catastrophic outcome would occur is unknown. However, even for well-studied problems like global warming, we know very little about the likelihood of catastrophic impacts. Furthermore, most other environmental problems (e.g., deforestation and toxic waste disposal) have received much less attention, so we know even less about the characteristics of the cost and benefit functions in these cases.

We have made greater progress in terms of understanding the discount rates that should be used to evaluate policies for which impacts occur over very long time horizons. The very fact that the marginal return on capital, and thus real interest rates, are stochastic implies that the effective discount rate should be lower than some kind of average expected rate, and a rate close to zero is not implausible. But what matters is the tail of the distribution, i.e., the probabilities of one or more severe economic contractions over the next century, and given our (fortunately) limited experience with severe contractions, we are left with ranges of plausible discount rates. And while zero and two percent may both be plausible numbers, they can have very different implications when policy benefits occur 100 years from now.

The good news is that environmental economists have plenty of work to do, and need not plan on early retirement. Likewise for physical scientists working on environmental problems. It seems to me that at least one major focus of research should be on the causes and likelihood of severe or catastrophic outcomes, and this will likely involve collaboration between economists and physical scientists. And while global warming is an important problem, more work is needed on other problems that may eventually turn out to be even more pressing, such as the depletion or degradation of water resources, acid rain, toxic (and non-toxic) waste disposal, and the loss of wilderness and wildlife.

References

- Arrow, Kenneth J., and Anthony C. Fisher, “Environmental Preservation, Uncertainty, and Irreversibility,” *Quarterly Journal of Economics*, 1974, **88**, 312—319.
- Barro, Robert J., “Rare Disasters and Asset Markets in the Twentieth Century,” unpublished manuscript, December 2005.
- Bernanke, Ben S., “Irreversibility, Uncertainty, and Cyclical Investment,” *Quarterly Journal of Economics*, 1983, **98**, 85—106.
- Chakravorty, Ujjayant, James Roumasset, and Keping Tse, “Endogenous Substitution among Energy Resources and Global Warming,” *Journal of Political Economy*, 1997, **105**, 1201—1234.
- Clark, Colin W., *Mathematical Bioeconomics*, 2nd Ed., New York: John Wiley, 1990.
- Clarke, Harry R., and William J. Reed, “Consumption/Pollution Tradeoffs in an Environment Vulnerable to Catastrophic Collapse,” *Journal of Economic Dynamics and Control*, 1994, **18**, 991—1010.
- Conrad, Jon, “Quasi-Option Value and the Expected Value of Information,” *Quarterly Journal of Economics*, 1980, **92**, 813—819.
- Dixit, Avinash K., and Robert S. Pindyck, *Investment Under Uncertainty*, Princeton University Press, 1994.
- Dybvig, Philip H., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, “Long Forward and Zero-Coupon Rates Can Never Fall,” *Journal of Business*, 1996, **69**, 1—25.
- Fisher, Anthony C., and W.M. Hanemann, “Information and the Dynamics of Environmental Protection: The Concept of the Critical Period,” *Scandinavian Journal of Economics*, 1990, **92**, 399—414.
- Fisher, Anthony C., and Urvashi Narain, “Global Warming, Endogenous Risk, and Irreversibility,” *Environmental and Resource Economics*, 2003, **25**, 395—416.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, June 2002, **40**, 351—401.
- Freeman, A. Myrick, “The Quasi-Option Value of Irreversible Development,” *Journal of Environmental Economics and Management*, 1984, **11**, 292—294.

- Gollier, Christian, "Discounting an Uncertain Future," *Journal of Public Economics*, 2002, **85**, 149—166.
- Gollier, Christian, *The Economics of Risk and Time*, Cambridge: MIT Press, 2001.
- Gollier, Christian, Bruno Jullien, and Nicolas Treich, "Scientific Progress and Irreversibility: An Economic Interpretation of the 'Precautionary Principle'," *Journal of Public Economics*, 2000, **75**, 229—253.
- Gore, Al, *An Inconvenient Truth*, Emmaus PA: Rodale Books, 2006.
- Goulder, Lawrence H., and William A. Pizer, "The Economics of Climate Change," NBER Working Paper No. 11923, January 2006.
- Hanemann, W. Michael, "Information and the Concept of Option Value," *Journal of Environmental Economics and Management*, 1989, **16**, 23—37.
- Heal, Geoffrey, and Bengt Kriström, "Uncertainty and Climate Change," *Environmental and Resource Economics*, June 2002, **22**, 3—39.
- Henry, Claude, "Investment Decisions Under Uncertainty: The Irreversibility Effect," *American Economic Review*, Dec. 1974, **64**, 1006—1012.
- Intergovernmental Panel on Climate Change (IPCC), "Climate Change 2001: The Scientific Basis," Technical Report, 2001.
- Jacoby, Henry D., and A. Denny Ellerman, "The Safety Valve and Climate Policy," *Energy Policy*, 2004, **32**, 481—91.
- Kelly, David I., and Charles D. Kolstad, "Bayesian Learning, Growth, and Pollution," *Journal of Economic Dynamics and Control*, 1999, **23**, 491—518.
- Kelly, David I., and Charles D. Kolstad, "Malthus and Climate Change: Betting on a Stable Population," *Journal of Environmental Economics and Management*, 2001, **41**, 135—161.
- Kolstad, Charles D., "Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions," *Journal of Environmental Economics and Management*, 1996a, **31**, 1—18.
- Kolstad, Charles D., "Fundamental Irreversibilities in Stock Externalities," *Journal of Public Economics*, 1996b, **60**, 221—233.
- Mendelsohn, Robert, William Nordhaus, and D. Shaw. "Measuring the Impact of Global Warming on Agriculture," *American Economic Review*, 1994, **84**, 753—771. .

- Newell, Richard G, and William A. Pizer, “Discounting the Distant Future: How Much Do Uncertain Rates Increase Valuations?” *Journal of Environmental Economics and Management*, 2003a, **46**, 52—71.
- Newell, Richard G, and William A. Pizer, “Regulating Stock Externalities Under Uncertainty,” *Journal of Environmental Economics and Management*, 2003b, **46**, 416—32.
- Nordhaus, William D., *Managing the Global Commons*, Cambridge: MIT Press, 1994a.
- Nordhaus, William D., “Expert Opinion on Climate Change,” *American Scientist*, 1994b, **82**, 45—51.
- Pindyck, Robert S., “Uncertainty in the Theory of Renewable Resource Markets,” *Review of Economic Studies*, April 1984, **51**, 289—303.
- Pindyck, Robert S., “Sunk Costs and Sunk Benefits in Environmental Policy,” MIT Center for Energy and Environmental Policy Research, Working Paper No. 95-003, December 1998.
- Pindyck, Robert S., “Irreversibilities and the Timing of Environmental Policy,” *Resource and Energy Economics*, July 2000, **22**, 233—259.
- Pindyck, Robert S. “Optimal Timing Problems in Environmental Economics,” *Journal of Economic Dynamics and Control*, July 2002, **26**, 1677—1697.
- Pizer, William A., “Optimal Choice of Policy Instrument and Stringency Under Uncertainty: The Case of Climate Change,” *Resource and Energy Economics*, 1999, **21**, 255—87.
- Pizer, William A., “Combining Price and Quantity Controls to Mitigate Global Climate Change,” *Journal of Public Economics*, 2002, **85**, 409—434.
- Pizer, William A., “Climate Change Catastrophes,” Resources for the Future Working Paper 03-31, May 2003.
- Pizer, William A., “Climate Policy Design under Uncertainty,” Resources for the Future Working Paper 05-44, October 2005.
- Posner, Richard A., *Catastrophe*, New York: Oxford University Press, 2004.
- Roberts, Marc J., and A. Michael Spence, “Effluent Charges and Licenses Under Uncertainty,” *Journal of Public Economics*, 1976, **5**, 193—208.
- Roughgarden, Tim, and Stephen H. Schneider, “Climate Change Policy: Quantifying Uncertainties for Damages and Optimal Carbon Taxes,” *Energy Policy*, 1999, **27**, 415—429.

- Scheinkman, José A., and Thaleia Zariphopoulou, “Optimal Environmental Management in the Presence of Irreversibilities,” *Journal of Economic Theory*, 2001, **96**, 180—207.
- Singh, Rajesh, Quinn Weninger, and Matthew Doyle, “Fisheries Management with Stock Growth Uncertainty and Costly Capital Adjustment,” *Journal of Environmental Economics and Management*, Sept. 2006, **52**, 582—599.
- Stavins, Robert N., “Correlated Uncertainty and Policy Instrument Choice,” *Journal of Environmental Economics and Management*, **30**, 1996, 218—232.
- Tietenberg, Tom, *Environmental and Natural Resource Economics*, 7th Ed., Boston: Addison Wesley, 2006.
- Tsur, Yacov, and Amos Zemel, “Accounting for Global Warming Risks: Resource Management Under Event Uncertainty,” *Journal of Economic Dynamics and Control*, July 1996, **20**, 1289—1305.
- Ulph, Alistair, and David Ulph, “Global Warming, Irreversibility and Learning,” *The Economic Journal*, May 1997, **107**, 636—650.
- Weitzman, Martin L., “Prices vs. Quantities,” *Review of Economic Studies*, 1974, **41**, 477—491.
- Weitzman, Martin L., “Optimal Rewards for Economic Regulation,” *American Economic Review*, **68**, 683—691.
- Weitzman, Martin L., “Why the Far-Distant Future Should Be Discounted at its Lowest Possible Rate,” *Journal of Environmental Economics and Management*, 1998, **36**, 201—208.
- Weitzman, Martin L., “Landing Fees vs Harvest Quotas with Uncertain Fish Stocks,” *Journal of Environmental Economics and Management*, 2002, **43**, 325—338.
- Weitzman, Martin L., “Gamma Discounting,” *American Economic Review*, March 2001, **91**, 260—271.
- Weitzman, Martin L., “Prior-Sensitive Expectations and Asset-Return Puzzles,” unpublished manuscript, November 2006.