Uncertainty, Irreversibility, and the Timing of Climate Policy

Anthony C. Fisher
Department of Agricultural and Resource Economics
University of California at Berkeley

Prepared for the conference on the "Timing of Climate Change Policies"
Pew Center on Global Climate Change
October 2001

Contents

Acknowledgements

The Pew Center and author would like to thank Ev Ehrlich, C. D. Kolstad, and Alan Sanstad for their review and comments on previous drafts of this report.

Executive Summary

Much has been written in recent years about policy options for dealing with anticipated climate change, often, as in discussions of the Kyoto Protocol, with a focus on mechanisms to coordinate international efforts to control emissions of greenhouse gases. Another line of research, and the one pursued in this paper, considers the question of the timing of climate policy, or, in other words, how soon and how stringently to control emissions, in the face of a variety of uncertainties and irreversibilities, both physical and economic. Policy decisions (including a decision to do nothing, which is at one end of the spectrum of possible responses to climate change) must be made in the absence of a sure knowledge of the rate and degree of warming over the next several decades or centuries. Decisions must also be made under uncertainty about the physical impacts of a given degree of warming and the associated economic losses (or gains).

Further compounding the decision problem are rigidities or irreversibilities in both physical and economic systems. Investment in reducing emissions can be irreversible in the sense that capital embodied in nuclear or renewable energy production facilities, energy-efficient buildings, and so on, is sunk, i.e., cannot be readily converted to other uses in the event this is desired before the capital has fully depreciated. On the other hand, increases in atmospheric concentrations of greenhouse gases are, if not irreversible (there is depreciation or decay here too), very long lived, on human time scales, on the order of hundreds and even thousands of years. Further, impacts of an increase can be irreversible, and, indeed, catastrophic, as would be true for substantial melting of the Antarctic ice sheets and the attendant rise in sea level.

The problem considered in the paper is then to determine the implications of these kinds of uncertainties and irreversibilities for the timing of climate policy. The first section briefly reviews the environmental economics literature on decisions under uncertainty and irreversibility. The key finding, here, is that, where a problem is characterized by (1) uncertainty about future costs and benefits of the alternatives under consideration, (2) prospects for resolving or reducing the uncertainty with the passage of time, and (3) irreversibility of one of the alternatives, an extra value, an option value, properly attaches to the reversible alternative(s). Put differently, an irreversible decision or action has to clear a higher hurdle to pass a benefit/cost test. Just having a benefit/cost ratio greater than one, or a positive net present value, will not suffice. An illustrative empirical application, to the choice among alternative uses of a large tract of tropical forestland, suggests the potential importance of appropriately factoring relevant uncertainties and irreversibilities into the analytical framework.

The second section returns the focus to the climate problem and considers whether and how the findings from the more general environmental economics literature apply. The climate problem does appear to fit the conditions of what might be called the option-value theorem, but there is a difficulty: two sources of irreversibility that cut in different directions. Before the implied conflict is addressed, the climate irreversibility is discussed in some detail on the grounds that this is less familiar to economists and policy analysts than the investment irreversibility. Consideration of the long run necessarily raises questions of how to value costs and benefits that will accrue perhaps far in the future. Therefore, this section also reviews recent thinking on discounting in the context of global warming and other environmental problems similarly characterized by very long-lived or irreversible consequences. The key result, here, is that there is support, from several different approaches to discounting, for declining discount rates over time—declining perhaps to zero beyond some point several hundred years in the future.

The second section then goes on to review contributions to the literature on the timing of climate policy under uncertainty and irreversibility. The literature is rather sparse. Although voluminous on climate policy, and only a little less so on uncertainty and irreversibility generally, it has only just begun to focus on the implications of the latter for the former. A brief review suggests that the climate irreversibility is downplayed relative to the investment irreversibility, and the investment irreversibility cuts in the direction of slowing investment in reducing greenhouse gas emissions rather than speeding or increasing it. There appears to be relatively little cost, in the form of future damage from current emissions, to holding off on investment while waiting to learn what damage will occur.

Finally, the third section describes a model that precedes from somewhat different, arguably more realistic, assumptions about the climate irreversibility to somewhat different implications for the timing of climate policy. The key difference in assumptions is that emissions in a given period—the first period, in a two-period model—lead to an increase in the probability of a catastrophic impact in the next period. The models reviewed in the second section, with one exception, do not consider the possibility of a catastrophic impact, or “jump” in damages to some higher level. The exception considers the possibility of such a jump, but the probability that it will occur is not affected by first-period emissions and the resulting concentration of greenhouse gases in the atmosphere. When the probability is (positively) affected, the effect of the investment irreversibility is weakened. The tendency to decrease first-period investment is offset by the need to increase investment to reduce the probability of catastrophic second-period impact. Further, the slower the rate of decay of greenhouse gases in the atmosphere, the more first-period investment in reducing emissions is warranted.

Introduction

A great deal has been written over the past decade about policy options for dealing with anticipated climate change, including a range of theoretical and applied studies by academic economists. Many of these studies address the questions of what is an appropriate mechanism for coordinating international efforts to control emissions of greenhouse gases and what are the prospects for achieving control through negotiation or bargaining (see, for example, Hoel, 1997). Another line of research, and the one pursued in this paper, focuses on the questions of how soon, and how stringently, to control emissions. Often—though not explicitly in this paper—this takes the form of an analysis of the optimal level and trajectory of carbon taxes (see, for example, Peck and Teisberg, 1992).

Although there are exceptions, discussed below, it seems fair to say that, for the most part, the implications for policy of uncertainty about future climate changes, and rigidities or irreversibilities in physical and economic systems, have not been adequately analyzed. What is the nature of the uncertainty? First, it is uncertain how much global-mean temperature will increase over, say, the next 100 years—to say nothing of the more distant future. Second, impacts of a given increase are not definitely known. According to climate scientists, there is some probability, perhaps small, of catastrophic impact. For example, disintegration of the West Antarctic Ice Sheet would result in a rise in sea level of 5–6 meters. Recent research seems to suggest that more rapid (than anticipated) melting is going on in some parts of the ice shelf supporting the ice sheet, but it is not clear what is happening overall (Kerr, 1998). Although a “meltdown” is by no means certain, and may even be unlikely, at least over the next several hundred years, there is a nonnegligible probability that it will occur.

Another factor complicating a policy choice is the presence of rigidities, or irreversibilities, in both physical and economic systems. The irreversibilities in the context of climate policy are of two kinds. Investment in controlling emissions, for example, by accumulating a stock of capital in nuclear or renewable energy production, or more energy-efficient buildings, cars, and appliances, is irreversible in the sense that the capital is "sunk." Once the investment has been made, the resulting stock of capital is ordinarily not easily converted to consumption or other forms of capital, should one or another of these changes be desired at some future date, when the durable capital has not yet completely depreciated. On the climate side, the accumulation of greenhouse gases is more or less irreversible. Of course, there is "depreciation" here too as carbon dioxide is removed from the atmosphere by terrestrial vegetation or by the oceans. But this is a long process, and some fraction of emissions will in fact remain in the atmosphere for 1,000 years and beyond (Maier-Reimer and Hasselman, 1987; Shultz and Kasting, 1997; Joos, Muller-Furstenberger, and Stephan, 1999). Further, impacts of even a relatively short-lived pulse may be long-lived or irreversible. This would be true of the catastrophic impacts, such as disintegration of the West Antarctic Ice Sheet, or regional climate changes that resulted in significant loss of biodiversity.

This paper will attempt to answer the question: What are the implications for the timing of climate policy of various scenarios involving combinations of these uncertainties and irreversibilities? That is, the focus is on the question posed earlier, of how soon, and how stringently, to control greenhouse gas emissions, but with special reference to the implications of relevant uncertainties and irreversibilities.

The next section discusses uncertainty and irreversibility more generally in the context of environmental problems. Under what circumstances are they significant, and what are the implications for policy? The discussion, here, leads to the concept of option value, which, as will be shown, arises in a setting characterized by both irreversibility and the prospect of resolving or reducing uncertainty with the passage of time. This option value is exactly analogous to the value of a financial option but, as will be indicated, has a history of independent development in environmental economics. An application, to the problem of choosing among alternative uses of a large tract of tropical forestland, suggests the potential importance of accounting for option value in practice in decisions about the environment. The second section returns the focus to the climate problem, specifically, the timing of climate policy under uncertainty and irreversibility, with a review of relevant results from the rather sparse literature here. Since the review suggests that not all questions regarding the timing of climate policy are adequately addressed in the literature, the third sketches in nontechnical terms a simple two-period model of the optimal level of first-period investment in control of greenhouse gas emissions that in my judgment takes appropriate account of relevant uncertainties and irreversibilities. Results are compared and contrasted to those in the existing literature. A summary and conclusions are given in the fourth section.

I. Uncertainty, Irreversibility, and Option Value in Environmental Economics

The basic concepts and results concerning uncertainty, irreversibility, and option value in environmental economics are readily described. Uncertainty, is of course, a well-understood fact of economic life and often plays a role in decisions about the allocation of resources, including those of the natural environment. Economic models usually assume that the decision-maker does the best she can, subject to what is known, or expected, about future costs and benefits of any given course of action. An interesting question arises, however, concerning timing. Clearly, if a decision, say, about whether to go ahead with a development project in a natural environment, has to be made today, it has to be based on the expected values of benefits and costs, perhaps adjusted or discounted to reflect the risk preferences of the decision-maker or of those in whose name the decision is made. But suppose that the decision can be deferred, say, till next year. Why would one want to do this? If additional information about future costs and benefits will be forthcoming, perhaps resolving the uncertainty, presumably a better decision can be made—one that reflects a more accurate view of the advantages and disadvantages of going forward. Of course, deferring the decision comes at a cost: the foregone returns from the project over the time the decision has been deferred, which needs to be balanced against the benefit of the improved decision.

In a problem of this type, the option to postpone the irreversible investment decision until better information about future costs and benefits is available, rather than being constrained to make the decision right away, has a value: option value. This is the basic result of Arrow and Fisher (1974) and Henry (1974) in the environmental economics literature. Notice that the option has value only because the decision-maker is assumed to learn about future returns by waiting. If this were not the case, nothing would be gained by postponing a decision on whether or not to go ahead with the project. Option value has also been identified with the value of information (Conrad, 1980). This is correct but only in a qualified sense. The information about future costs and benefits is valuable only because the decision-maker is assumed to have the flexibility to postpone the decision about the investment. Option value can thus be considered a conditional value of information—the value of information conditional on retaining the option to either make the investment or not make it (Hanemann, 1989). It is worth noting also the parallel to finance. The option to postpone the investment decision is analogous to a call option on a share of stock. It confers the right to exercise the option to invest at a given price (the cost of the investment) to receive an asset (the project, say, a dam for the production of hydroelectric power) that will yield a stream of uncertain future returns.[1]

A. The Concept of Irreversibility

Although the decision problem, and the result, are easily stated, a good deal of structure is concealed or implicit in the statement. The problem will now be considered in more detail, focusing specifically on the concept of irreversibility, which, at least to economists, is less familiar and, perhaps, less justified, than the assumptions of uncertainty and the temporal resolution of uncertainty. Environmental economists are, perhaps, more sensitive than are other economists to an idea that runs through much of the noneconomic literature on the environment: that threatened environmental losses are significant because they will be experienced in perpetuity. This is not to say that all environmental losses are irreversible, but many of the major issues or conflicts seem to be about those that are: the loss of a wild river and canyon environment, such as the Grand Canyon, to water resource development; the loss of endangered species of plants and animals, or more generally, of biodiversity; and, of course, some of the potentially catastrophic impacts of global climate change.

When the proposition that irreversibility matters was put forward in the 1970s, in the work of John Krutilla and his associates at Resources for the Future (see, for example, Krutilla and Fisher, 1975; Fisher, Krutilla, and Cicchetti, 1972), it was met with somewhat contradictory responses on the part of academic economists: Everything is irreversible, in the sense that time does not run backwards; and nothing is irreversible, in that the consequences of any decision, for example, to develop a natural environment, can be reversed given sufficient inputs of conventional resources. In the unlikely event that the decision is not technically reversible, it should at least be economically reversible, in the sense that other goods or resources might be found to substitute in consumption for the lost natural environment. This latter response, that irreversibility is an “empty box,” was, perhaps, the dominant view among economists.

A rejoinder, to which most environmental economists today at least would probably subscribe, goes along the following lines. Consider the decision to develop a water storage reservoir in a more or less pristine environment. Correcting an ill-advised decision to build a dam involves more than simply dismantling the structure when the environmental costs are perceived to exceed the benefits. Supersaturation of the reservoir banks at full-pool elevations may result in sloughing and landslides into the reservoir during drawdown. Moreover, if streams of high turbidity are impounded, sediment will build. Dismantling the structure would then leave the impoundment area with an abiotic base quite different from that which originally existed. Perhaps this explains the strength of the opposition to the proposed damming of the Colorado River in the Grand Canyon, since the first proposals to do this early in the last century. At stake, in the view of the opponents, was and is permanent loss of two billion years of natural history recorded on the Canyon walls. Opposition is not based solely on the technical infeasibility of restoring the preproject environment. Also at issue are the preferences of individuals regarding the attributes of the environment. For some, authenticity in a natural environment is a valued attribute as it is to others in a work of art. A restoration or a copy, no matter how skillful, may not be a good substitute in consumption for the original article or environment. In that case a decision to construct a dam in the Grand Canyon would, indeed, be irreversible. Here, it is important to note that just because a decision is irreversible does not mean it should never be taken. The point is that irreversibility is a real phenomenon, and it does matter, in the sense that it does change the benefit/cost calculus, but this still involves balancing at the margin.

The dam construction example may be a bit dated. Today, the major environmental issues, at least on a global scale, are probably loss of biodiversity and climate change. Here, too, irreversibility seems to be significant in motivating concern. With respect to biodiversity, much of the concern is for endangered species. But, even if the survival of a particular species is not at issue, biological impacts can be very difficult to reverse over any time span that is meaningful for human societies. The clear-cutting of a climax forest species, for example, removes the results of an ecological succession that may represent centuries of natural processes. Further regeneration may not lead to the original configuration even after many more centuries. Opportunistic species, such as hardy grasses, may come in and preempt the niche otherwise filled, eventually, by the climax species (Albers and Goldbach, 2000). With respect to climate change, relevant rigidities or irreversibilities have been noted in the introduction and are discussed at greater length in the second section.

An illustrative empirical application, in Box 1 below, provides a powerful demonstration of the importance of the concepts discussed in this section, specifically, of appropriately accounting for uncertainty and irreversibility in environmental policy.

Box 1: An Illustrative (Non-Climate) Application

The application, taken from Albers, Fisher, and Hanemann (1996), is to valuation and management of a tract of forest land in Thailand, specifically, to the problem of allocating the land among three competing uses, over three periods (defined by the timber rotation time), subject to constraints on the conversion from one use to another. Thus, it is feasible to go from activities compatible with preservation (including hunting and gathering, ecotourism, and erosion control) to "intermediate" activities, such as small-scale shifting cultivation and extraction, and from either of these kinds of uses to development, in this case eucalyptus plantations and permanent agriculture. It is not feasible to go in the other direction, for example, from commercial agriculture to "preservation."

The area studied is partly included in an existing national park, Khao Yai National Park, in central Thailand. The analysis divides the area into four management units or plots. The outer edge of the Park, plot 1, has been encroached and begins in the intermediate use category. The inner plots, 2 and 3, begin in the preservation category. Plot 4, also in preservation, is not currently in the Park but is being considered for inclusion. On the basis of estimates of the benefits of the alternative uses, and assumptions about the probability of different outcomes (for example, rapid growth of tourism revenues versus slow or no growth), a programming model determines the pattern of allocation of each of the plots in each period that maximizes the present value of the entire area over all three periods. This is done in two ways. First, it is assumed that no new information about future benefits will be forthcoming, or that the prospect of new information is ignored, in the first-period decision. Future benefits are simply replaced by their expected values as of the first period. This leads to a standard present-value benefit-cost analysis, which in turn calls for converting plots 1 and 3 to development, preserving plot 2, and converting plot 4 to intermediate uses.

The second approach assumes that there will be new information about future benefits, specifically, that uncertainty about the benefits of each of the alternative uses in each future period will be resolved by the beginning of the period and that this is recognized in the first-period decision. In this case, the optimal allocation calls for leaving plot 3 in preservation with the other plots used in the same way as in the first case. Since plot 3, in the interior of the present Park, however, accounts for about half of the total area, the difference is dramatic. This is true even though the calculated option value is just under 2 percent of the total value over the three periods. In this application, at least, it turns out that a relatively small option value, when appropriately measured and accounted for, can make a large difference in the outcome of an environmental decision. Here, it should also be noted that the option value is probably understated, as it reflects very conservative assumptions about the degree of uncertainty surrounding the alternative uses, and it has been shown elsewhere that option value will be larger in situations of large divergences in possible outcomes (Fisher and Hanemann, 1986).

B. A Brief Restatement of Theoretical Results

To review the discussion to this point of uncertainty, irreversibility, and option value in environmental economics, the main findings and implications for policy will be restated. Where a decision problem is characterized by (1) uncertainty about future costs and benefits of the alternatives, (2) prospects for resolving or reducing the uncertainty with the passage of time, and (3) irreversibility of one or more of the alternatives, an extra value, an option value, properly attaches to the reversible alternative(s). This is the value of retaining the option to choose any of the alternatives in the light of new information—an option that is necessarily lost if the irreversible alternative has already been chosen. Further, it appears that many environmental problems can be characterized in this way. The implication for policy is that an irreversible decision or activity has to clear a higher hurdle in order to pass a benefit/cost test. Just having a positive net present value, or a benefit/cost ratio greater than one, will not suffice. Here, it must be noted that calculation of the “correction factor” or option value in an empirical application is no easy task as it involves knowledge of relevant probability distributions of benefits and costs and how these distributions evolve over time. But the qualitative result, at least, that the standard net present-value or benefit/cost comparison is biased in favor of the irreversible alternative, is clear.[2]

II. The Climate Problem

Returning to the climate problem, the manner in which the concepts and results described in the preceding section apply will be considered. As noted in the introduction, the problem of how soon, and how stringently, to control greenhouse gas emissions is characterized by both uncertainties and irreversibilities. The uncertainties have to do with both the degree of warming to be expected and the potential impacts. Early and widely influential analyses of the role of uncertainty in climate/economy models, such as those of Manne and Richels (1999) and Nordhaus (1994), tend to conclude that uncertainty has only marginal effects on policies derived from deterministic versions of the models. These models do not however include irreversibilities. Further, they make a crucial assumption, as discussed below, that the probabilities of different levels of damages from warming are exogenous, i.e., are not affected by model variables such as atmospheric concentrations of greenhouse gases or the degree of warming.

The irreversibilities are on both the climate side and the investment (in emission control) side. Does the basic result on option value, or what might be called the "irreversibility effect" in environmental decisions, such as those discussed in the preceding section, provide any information about the timing of climate policy? Since it seems plausible to assume that additional information about the nature and impacts of climate change will be forthcoming over time, all of the conditions for the result to go through are in place, so the answer is yes—but. The "but" refers to the presence, in this problem, of two conflicting sources of irreversibility or inflexibility: on the one side, sunk capital (to reduce greenhouse gas emissions), and, on the other, a not-fully-degradable stock of greenhouse gases in the atmosphere and possibly irreversible impacts. The difficulty is that these cut in opposite directions. Therefore, it is not clear whether the conditions of the problem imply that investment in control ought to be slowed or reduced, while waiting for information needed to make a better decision, or that investment should come sooner to preserve the option to protect ourselves from impacts that may be revealed in the future as serious or even catastrophic.

It may be appropriate to characterize a debate that has sprung up over this dilemma as pitting mainly economists on the side of waiting to invest until better information is available and mainly natural scientists on the side of doing more now to forestall what they see as potentially serious impacts. Of course, this is something of a stereotype and does not fit everyone. In the remainder of this section, the concept of irreversibility as it applies to the climate problem is considered and some of the major contributions by economists who have developed models of the timing of climate policy that explicitly take account of uncertainty and irreversibility are reviewed. In the next section a new model that takes better account of these phenomena will be described and how this affects results and implications for policy will be indicated.

A. Irreversibility and Climate Change

*It seems useful to say a bit more about irreversibility in this context, as the relevant natural science is, perhaps, not well known to economists and thus not yet adequately incorporated in economic models that deal with climate change***.** In the introduction it is acknowledged that the accumulation of greenhouse gases in the atmosphere is not truly irreversible since, though emissions in any period cannot be negative, the resulting accumulation is subject to natural decay over time. Typically, in economic models, such as the well-known and widely-used DICE (Dynamic Integrated Climate/Economy) model, developed by Nordhaus (1993, 1994), the process of accumulation and decay is represented in a single equation, in which the accumulation, or stock, of CO_2, the main greenhouse gas, in a period (a decade in DICE) is equal to some fraction of emissions in the preceding period plus some other fraction (one minus the rate of decay over the decade) of the stock in the preceding period. The process continues unchanged over time, implying that the atmospheric concentration of CO_2 returns to its current level in about 300 hundred years and to the preindustrial level within 1,000 years. Yet, as at least a couple of recent contributions by natural scientists have pointed out, this is not likely. The difficulty is that, after relatively rapid mixing, over a few decades, of the atmosphere with the surface ocean, further removal of CO_2 from the atmosphere depends on mixing of the surface ocean with the deep ocean—a much slower process (Joos, Muller-Furstenberger, and Stephan, 1999). According to one calculation, after 1,000 years, CO_2 concentrations will still be well over twice the current level and nearly three times the preindustrial level and will remain elevated for several thousand years (Schultz and Kasting, 1997).

The point about long-run concentrations in fact appeared in the climate science literature over a decade earlier (Maier-Reimer and Hasselman, 1987) and is incorporated in at least a couple of insightful and important—but thus far largely neglected—economic analyses, by Azar and Sterner (1996), who calculate the impact on estimates of the damages from global warming, and by Farzin and Tahvonen (1996), who derive the implications for the optimal trajectory of carbon taxes. The most recent version of the DICE model (Nordhaus and Boyer, 2000) also takes account of the criticism of the representation of the carbon cycle in the earlier version. The one-equation model of carbon decay in the atmosphere is replaced by a three-equation, three-medium, model designed to represent the relatively slow exchange between the surface or upper ocean layer and the deep ocean. It turns out that this makes very little difference to the simulated optimal trajectory of rates of emissions control, but, most likely, this result is in large part due to the way in which damages beyond the next few decades are discounted, as discussed below.

Clearly, long-run concentrations—and also long-run or irreversible impacts, such as would result from the large scale melting of the Antarctic ice sheets—only matter if the discount rate used in evaluating a program of emissions control is sufficiently low that damages occurring or persisting several hundred and, indeed, several thousand years in the future matter to those who will be making the decisions about investment in control today. Azar and Sterner show that a change in the earlier version of the DICE model representation of the carbon cycle, so that it more accurately reflects long-run atmospheric concentrations, as in Maier-Reimer and Hasselman, does not make much difference in the optimal rate of emission control. This is because long-run damages are so heavily discounted—just as in the new version of the DICE model with the modified carbon cycle. In turn, this raises the question of what is the appropriate discount rate (or, more generally, discounting procedure) for dealing with environmental problems, such as global warming, or management of the nuclear fuel cycle, characterized by very long-lived consequences. Interestingly, it is the consideration of just such problems that has stimulated recent work on whether and how conventional discounting needs to be modified to deal adequately with the relatively distant time horizons. The arithmetic of conventional exponential discounting is inexorable and leads to the insignificance of virtually any impact beyond 50–100 years in the future. Yet, most natural scientists, perhaps the general public, and even economists who consider problems, such as global warming, with its potentially catastrophic consequences emerging only in the more distant future, have a sense that the conventional approach to valuation in these circumstances is somehow lacking.

A full discussion of the choice of an appropriate discount rate, or, more generally, discounting procedure, is beyond the scope of this paper, but the tendency of recent thinking about discounting in the context of global warming is briefly reviewed in Box 2 below. The conclusion is that an appropriate social discount rate for evaluating future consequences of global warming should probably be declining over time, perhaps to zero beyond some point. This in turn suggests that long-run impacts matter, at least more than they would if simply discounted in the conventional way.

Box 2: A Digression on Discounting

The social discount rate is ordinarily given, in dynamic models, such as growth models in the tradition of Ramsey (1928), by the sum of the rate of pure time preference, or the utility discount rate, and the growth discount rate, where the latter is given by the elasticity of the marginal utility of consumption times the rate of change in consumption. The second term, the growth discount rate, reflects the assumption that a given increment in consumption will be worth less in the future when everyone is richer. This is the formulation also in the DICE model, which is a growth model joined to a climate model. Note that, even if the rate of pure time preference were zero, which some, including Ramsey, have argued it should be, to avoid favoring any one period or generation relative to any other, the growth discount rate, and, therefore, the social discount rate, would be positive, as long as consumption is expected to increase over time. Suppose, for example, that consumption is expected to increase by 3 percent from one period to the next. If the elasticity of marginal utility is one, i.e., if marginal utility declines by x percent when consumption increases by x percent, then the growth discount rate in the example will be 3 percent, which would also be the social rate of discount if the utility discount rate were zero.

Another important point about this formulation is that it implies that, as the rate of growth in consumption changes, the discount rate changes. In particular, if the rate of growth in consumption declines, so does the discount rate. This is reflected in the DICE simulations, in which total factor productivity is assumed to be increasing, but at a decreasing rate, yielding a social discount rate that decreases from 5.9 percent in 1995 to 4.4 percent in 2075 (Nordhaus, 1994, p. 91). Since the utility discount rate is, however, assumed to be 3 percent, this is a lower bound for the social discount rate. Azar and Sterner put forward a number of arguments in favor of a zero or near-zero social rate of time preference even though market data may suggest a nonzero private rate. Both they and Dasgupta (2000) are also considerably more pessimistic than Nordhaus about prospects for continued growth in consumption at anything like the rates observed over the past few decades, especially when prospects for long-term environmental degradation, and loss of ecosystem services, are taken into account. This last point is important. Suppose, as a result of global warming and, perhaps, other changes having an adverse impact on ecosystem services, consumption at some future date actually declines. Then the growth discount rate would become negative. If the utility discount rate were close to zero, the social rate of discount would also become negative. In this scenario, the value of future damages would be amplified, not reduced, by discounting. Alternatively, the return on an investment that produced a reduction in greenhouse gas emissions, and, therefore, in warming, would be increased, rather than decreased, by discounting.

It seems fair to say that, in the widely-accepted framework of growth models following Ramsey, there is a consensus that the discount rate used to evaluate damages from global warming, or investment in reducing damages, is likely to be changing over time, declining if consumption is increasing, though agreement does not extend to how fast, or how far, the discount rate declines or to where it starts. In the DICE model, the discount rate starts fairly high, and does not decline very much, in part because of the constant positive rate of time preference. Consequently, damages in the more distant future are given very little weight. Azar and Sterner (1996, p. 181) show that, even assuming continued, though declining, growth in consumption over a time horizon of 300 years, but alternative lower rates of time preference, the marginal cost of CO_2 emissions would rise from $13 per ton, at a time preference or utility discount rate of 3 percent, to $32 at 1 percent, to $75 at 0.1 percent, and finally to $85 at 0 percent.

A quite different approach to discounting in the context of the climate problem, rooted explicitly in uncertainty, is put forward by Weitzman (1998). Interestingly, it points in the same direction, namely, declining—in this case dramatically declining—discount rates over time. Weitzman begins by observing that there is at any time, say, the present, a wide distribution of beliefs, even among the "experts," about the social discount rate, based on differences in beliefs about rates of time preference, future growth, the impact of environmental degradation or other externalities, and so on. Representing the distribution of beliefs explicitly by a probability distribution, Weitzman shows that the implied social discount rate, an aggregation of the individual beliefs, declines over time because increasingly greater present-value weight is placed on stated or observed low discount rates. High discount rates are relatively less important over time because their present value is progressively reduced relative to that of the low discount rates. Eventually, the only rate that counts is the minimum rate.

In subsequent work Weitzman (2001) specifies a particular probability distribution for discount rates, the gamma distribution, characterized by a single peak, with the left tail cut off at zero and the right tail asymptotically approaching the x-axis, and shows it to be consistent with the results of a survey of the stated beliefs of a large sample of academic economists. The implied social discount rate starts off at the mean value of the distribution and declines monotonically towards zero. On the basis of the mean and variance of the sample, and a couple of other numerical assumptions, the discount rate starts off at 4 percent and remains there through years 1–5 (what Weitzman calls the "immediate future"), falls to 3 percent for years 6–25 (the "near future"), to 2 percent for years 26–75 (the "medium future"), to 1 percent for years 76–300 (the "distant future"), and finally to zero beyond 300 years (the "far-distant future").

B. The Literature on Climate Change, Uncertainty, and Irreversibility: Approaches and Results

Granted, then, that irreversibilities matter, what does the literature reveal about the implications for the timing of climate policy of the apparently conflicting or opposing irreversibilities? Surprisingly, perhaps, results here tend to downplay the importance of the climate irreversibilities discussed above, relative to the investment irreversibility. The opposing effects were first recognized and jointly analyzed by Kolstad (1996a) in a two-period model of irreversibilities in stock externalities. (A stock externality is an externality that, like some forms of pollution, accumulates—and, perhaps, decays—over time.) Kolstad asks the question, how does the prospect of better second-period information about the consequences of the externality, in his example the damages from global warming, affect the desired level of first-period investment in abatement capital? Emissions are assumed to be nonnegative, and the degree of capital "sunkness" and the decay rate of the stock of greenhouse gases are fixed. He finds that, if learning is proceeding slowly enough, compared with the rates of pollution decay and capital depreciation, learning makes no difference. If, on the other hand, learning is significant, either or both of the irreversibilities can affect the desired level of first-period emissions in opposite directions. Which dominates depends on the relative magnitudes of the decay and depreciation rates and on expectations about damages.

In a second paper, a multiperiod numerical simulation of optimal investment in control of greenhouse gas emissions based on the DICE model, here, extended to include both the capital-stock irreversibility and a parametric representation of the rate of learning, Kolstad (1996b) finds a significant impact associated with the capital-stock irreversibility but not with the emissions irreversibility. The reason, essentially, is that in his parameterization of the model, the nonnegativity restriction on emissions is never binding. Too little investment in emissions control in the early periods can be compensated by a bit more investment in later periods. There is no scenario in which it would be optimal to emit negatively in the future to correct for emitting too much today.

This is consistent with the main result in Ulph and Ulph (1997), a two-period model of global warming, irreversibility, and learning in which there is no explicit representation of investment in abatement, but, as in Kolstad, emissions are restricted to be nonnegative, the decay rate of the stock of greenhouse gases is fixed, and there is learning about damages. A sufficient condition for there to be an irreversibility effect, i.e., for first-period emissions with learning to be less than first-period emissions with no learning, is that the nonnegativity restriction is binding in the no-learning case. Ulph and Ulph also provide a multiperiod numerical simulation. For a variety of scenarios, they find very little difference between first-period emissions with learning and without. In one case, characterized by a low discount rate and substantial uncertainty, emissions with learning are greater. Since there is no explicit capital-stock irreversibility in the model or the simulation, it is not clear what is driving this somewhat counterintuitive result and the authors provide no explanation.

Results with much the same flavor as Kolstad's are obtained in a recent theoretical analysis of irreversibilities in environmental policy more generally (i.e., not just in climate policy though this certainly fits the framework) by Pindyck (2000). Pindyck's work also draws on concepts and methods in option-pricing theory as developed in the literature on finance. He represents the uncertainties about both the future costs and benefits of reduced environmental degradation, as could be accomplished by an investment in pollution control, and the future evolution of the environment itself, as particular stochastic processes (geometric brownian motions in which the percentage rate of change in each of the values is a random walk with drift, i.e., with a trend). This seems fairly plausible, since one characteristic of the process is that the variance increases with the passage of time, and it also accommodates a positive trend, a negative trend, or no trend. It turns out that the geometric Brownian motion is mathematically convenient and is, perhaps, for this reason as well, widely used in economic applications, including those bearing on the use of natural and environmental resources (see, for example, Brennan and Schwartz, 1985, and Conrad, 1997).

Pindyck too recognizes the opposing irreversibilities. He asks the question, given the opposing biases in a decision, say to invest in control of greenhouse gas emissions, is it possible to draw any general conclusions? Interestingly, the answer is "yes" though this does depend on some assumptions. One assumption is that the decision can be deferred; information about the evolving environmental impacts and values accrues over time and can lead to a better decision. Why make an irreversible investment in reducing emissions today if waiting a little while will reveal with greater certainty what the outcome will be? Waiting entails a cost, the foregone benefits from the investment during the waiting period, but this may be much less than the benefits from the better decision that results. Of course, the same argument applies on the other side: better to protect the environment from irreversible damage now since it can be "unprotected" if new information suggests that the damage, though irreversible, will be minor. What Pindyck shows is that an increase in uncertainty, whether over future costs and benefits of environmental protection or over the behavior of the environment, leads to a higher threshold for policy adoption. Policy adoption involves a sunk cost associated with a reduction in the entire trajectory of future emissions whereas waiting involves only continued emissions over the waiting period. As Pindyck notes, this result depends on the extent to which the policy is, indeed, irreversible. It is important to note also that it implicitly assumes that emissions over the period do not increase the risk of a catastrophic impact.

To sum up, the sense one gets from the rather sparse economic literature that considers the implications for the timing of climate policy of uncertainty and irreversibility is that these cut in the direction of slowing or reducing investment in control of greenhouse gas emissions rather than speeding or increasing it. The reason for this result, which must appear counterintuitive to many climate scientists and, perhaps, also to policy-makers and to the general public is that, given the assumptions or parameter values built into the models, there is relatively little cost to holding off on investment while waiting to learn of the benefit. The climate irreversibility, though recognized, does not play much of a role in driving current decisions on controlling emissions. In the next section a model derived from somewhat different assumptions that yields somewhat different results is described.

III. Climate Change, Irreversibility, and Endogenous Risk

In discussing Pindyck's model, just above, it was noted that a key (implicit) assumption driving his results is that emissions in a given time period do not lead to an increase in the probability of a catastrophic impact at some point in the future—or the next period, in a two-period model. The model described in this section, by contrast, explicitly assumes that the probability of such an event may be positively related to the level of greenhouse gas concentrations in the atmosphere, or in other words, that the probability is endogenous to the model. This appears more realistic, but the results obtained are easily compared to those that follow from the contrary assumption, namely, even if there is some possibility of catastrophic impact, the probability of this occurring is not explicitly related to any of the variables determined within the model, i.e., is exogenous. It is worth emphasizing, here, that nothing in the model says that a catastrophic impact is likely. It may be very unlikely. It is assumed only that there is a (possibly very small) nonzero probability that some such event will occur, and that the probability may be related to the degree of pollution, as measured by concentrations of greenhouse gases.

The key features of the model are: two discrete time periods (think of present and future); sunk or irreversible investment in controlling emissions; nondegradable or irreversible stocks of greenhouse gases; possibly endogenous risk of catastrophic damages; and future learning about the nature of damages (Fisher and Narain, 2001). Learning is fixed, in the sense that the decision-maker is assumed to learn, by the start of the second period, whether a climate event—say, a 5°F rise in global-mean temperature—has occurred and, if it has, the nature of the impact, high damage or low. The model then studies how the desired level of first-period investment varies with the degree of sunkness of the resulting capital stock, things such as energy-efficient buildings, renewable energy production facilities, nuclear plants, and so on. Also considered is how the desired level of investment varies with the degree of nondegradability of the stock of gases. Of course, in reality there are several different gases, with different residence times in the atmosphere, but these are aggregated to keep the model as simple as possible with no significant loss in generality.

One difference from Kolstad's model, here, is in the definition of sunk capital. Kolstad defines this in terms of durability whereas Fisher and Narain prefer to follow the literature on investment and define it in terms of convertibility: Capital is sunk if it cannot be converted into consumption or other forms of capital. As it turns out, results are unaffected by the definition though in a different setting it can make a difference (Narain and Fisher, 2000).

A few words about the interpretation of time scales in the context of a two-period model may be in order. The first period can be as short as 10 or 20 years if this is sufficient for observation or research to yield inferences about the degree of warming over the next several decades. In this case, capital will not be fully depreciated, and, if it is also sunk, in the sense that it cannot be easily or cheaply converted to consumption or other forms of capital, the decision-maker may, at the start of the second period, regret the first-period investment. If, on the other hand, it is appropriate to regard the first period as somewhat longer, say, 50 or 100 years, then, even if capital is sunk, there will be little to regret since it will have fully depreciated by the start of the second period. In such a case the investment irreversibility will not matter very much and a first-period decision on investment will be dominated by the climate irreversibility. Since previous studies have, however, emphasized the consequences of investment irreversibility, a model that at least allows for this possibility is preferred.

The structure of the model can be described quite simply. An economic agent is assumed to allocate a fixed resource endowment to either consumption or investment (in reducing emissions of greenhouse gases) in each of two periods. The agent also has the option of increasing consumption by disinvesting (at a cost), i.e., by converting some of the capital back into consumption. The point of the model will be to determine how the desired or optimal level of first-period investment varies with the ease or difficulty of disinvestment in the second period as well as with the persistence of atmospheric concentrations of greenhouse gases.

The formal objective is to choose investment and consumption in each period to maximize the sum of benefits, or utility, over both periods. Since the agent learns about second-period damages at the start of the second period, the optimization problem is solved through backwards induction. This entails first choosing the optimal level of investment in the second period when the returns are known. The second-period choice yields in the first period what is known as the expected continuation value, expected second-period returns given that second-period investment is optimally chosen. The problem then boils down to choosing the level of first-period investment given both first-period returns and the expected continuation value. A simple numerical example illustrating this analytical approach, known as stochastic dynamic programming, is given in an Appendix.

Results of the model can be summarized in the following propositions.

1. Investment in the first period is a decreasing function of the degree of sunkness of capital if risk is exogenous. In other words, the more difficult or expensive it is to "disinvest," the less investment there should be in the first period. This is the Kolstad-Pindyck result though arrived at somewhat differently.

2. The sign of the relationship between first-period investment and the degree of sunkness is ambiguous if risk is endogenous. The tendency to decrease investment as capital becomes more sunk is compensated by the need to increase investment to reduce the probability of a catastrophic impact in the case where the probability is related to the level of greenhouse gas concentrations.

3. Investment in the first period is a decreasing function of the rate of decay of the stock of greenhouse gases if risk is exogenous. The more persistent are greenhouse gases in the atmosphere, the more first-period investment in reducing emissions is warranted.

4. Investment in the first period is a decreasing function of the rate of decay of the stock of greenhouse gases if risk is endogenous under three sufficient conditions. In essence the conditions require that the decision-maker not have an incentive to trigger the climate event (say, a 5°F rise in global-mean temperature with associated probabilities of high and low levels of damages) and are quite plausible.

These results have a somewhat different flavor than those obtained by Kolstad, Ulph and Ulph, and Pindyck. Loosely speaking, Fisher and Narain find more of a tendency for the irreversibility associated with the accumulation of greenhouse gases in the atmosphere to matter and a weaker effect of irreversibility associated with investment in reducing emissions. No doubt the differences in results reflect differences in assumptions, especially the assumption that the risk of climate change is endogenous, i.e., is positively related to the size of the accumulated stock of greenhouse gases. This assumption appears to better reflect physical reality, but, whether it does or not, the model of this section demonstrates that implications for the timing of climate policy can be quite sensitive to the way in which that reality is reflected in economic models.

IV. Conclusions

This paper has considered the question of how the timing of climate policy (i.e., the timing of investment in reductions in greenhouse gas emissions) is affected by uncertainties and irreversibilities. The uncertainties are about both the degree of warming and potential impacts. Irreversibilities are manifested in both physical systems (emissions cannot be negative, impacts may be very long lasting or impossible to undo) and economic systems (capital embodied in nuclear or renewable energy production facilities or energy-efficient buildings is sunk, i.e., cannot be readily converted to other uses).

A review of the environmental economics literature (first section) yields a key result: Where a problem is characterized by (1) uncertainty about future costs and benefits of the alternatives under consideration, (2) prospects for resolving the uncertainty with the passage of time, and (3) irreversibility of one of the alternatives, an extra value, an option value, is attached to the reversible alternative(s). Put differently, an irreversible decision or action has to clear a higher hurdle to pass a benefit/cost test. An illustrative empirical application, to the choice among alternative uses of tropical forestland, suggests that explicitly accounting for uncertainty and irreversibility in this fashion can make a big difference.

Application to the climate problem (second section) is not straightforward since there are two opposing irreversibilities that cut in different directions. The sparse economics literature to date on this problem suggests that the investment irreversibility is more important than the climate irreversibility and the investment irreversibility cuts in the direction of slowing investment in reducing greenhouse gas emissions. As the models are formulated, there is relatively little cost, in the form of future damage from current emissions, to holding off on investment while waiting to learn what damage will occur.

A new model (third section) differs primarily in an assumption that the probability of high damages, or catastrophic impact, is positively related to the level or concentration of greenhouse gases in the atmosphere. Thus, emissions in an early period—the first period in a two-period model—lead to an increase in the probability of a catastrophic impact in a later, or second, period. In this case the effect of the investment irreversibility is weakened. The tendency to decrease early-period investment is offset by the need to increase investment to reduce the probability of late-period impact. Also, the slower the rate of decay of greenhouse gases in the atmosphere, the more early-period investment in reducing emissions is warranted.

Appendix A: Stochastic Dynamic Programming: a Simple Example

Suppose an investment decision is to be made now, on the basis of the cost, (assumed known) returns in the first period, and expected returns in the second. Suppose, further, that there are just two possible second-period outcomes—one good and one bad. Let the investment cost be I = 12, first-period returns b_0 = 10, the good second-period outcome $\overline{b}_1$ = 15, the bad second-period outcome $\underline{b}_1$ = 5, the probability of the good outcome p = 0.9, the probability of the bad outcome (1 – p) = 0.1, and the discount rate r = 0.1.

An expression for returns to the investment, over both periods, can be written as

$$V_0 = b_0 + \frac{p\overline{b}_1 + (1 - p)\underline{b}_1}{1 + r} .$$

Substituting the assumed values for $b_0, p, \overline{b}_1, \underline{b}_1$, and r,

$$V_0 = 10 + \frac{0.9(15) + 0.1(5)}{1.1} = 12.73$$

to the nearest hundredth. The expected net present value of the investment is thus

$$V_0 - I = 12.73 - 12 = 0.73.$$

The optimal choice in the first period is thus to go ahead and make the investment. The investment criterion can be written more formally: defining W_0 as the net payoff,

$$W_0 = \max\{V_0 - I, 0\}.$$

That is, the net payoff is either $V_0 - I$, if the investment is made, or 0, if it is not, and it will pay to invest as long as $V_0 - I > 0$, which it is in this case.

Now, suppose the investment opportunity will be available in the second period if it is not taken in the first. Returns to the investment undertaken in the second period are

$$V_1 \begin{cases} = \overline{b}_1 & \text{if } b_1 = \overline{b}_1 \\ = \underline{b}_1 & \text{if } b_1 = \underline{b}_1 \end{cases}.$$

The net payoff to an optimal decision in the second period, called the continuation value, is

$$F_1 = \max\left\{V_1 - I, 0\right\}.$$

What is the implication for the first-period decision? Notice that, although the second-period decision is made on the basis of known costs and returns (by the second period, the decision-maker knows whether $b_1 = \overline{b}_1$ or $\underline{b}_1$ and chooses accordingly), from the perspective of the first period, both V_1 and F_1 are uncertain. From the perspective of the first period, the expected continuation value is

$$\begin{aligned} E_0\left[F_1\right] &= p \max\left\{\overline{b}_1 - I, 0\right\} + \left(1 - p\right)\max\left\{\underline{b}_1 - I, 0\right\} \\ &= 0.9 \max\left\{15 - 12, 0\right\} + 0.1 \max\left\{5 - 12, 0\right\} \\ &= 0.9\left(3\right) + 0.1\left(0\right) \\ &= 2.7. \end{aligned}$$

The net payoff to the investment opportunity presented in the first period, recognizing that a decision can be postponed to the second period, is

$$\begin{aligned} F_0 &= \max\left\{V_0 - I, \frac{E_0\left[F_1\right]}{1 + r}\right\} \\ &= \max\left\{12.73 - 12, \frac{2.7}{1.1}\right\} \\ &= \max\left\{0.73, 2.45\right\}. \end{aligned}$$

The optimal decision in the first period is now to hold off on the investment since the expected continuation value (2.45) exceeds the net present value of the investment made in the first period (0.73). The difference between F_0 and W_0, in this example $2.45 - 0.73$, can be interpreted as option value, since F_0 is the maximum value of the investment opportunity in the first period taking into account the option to defer the decision until the second period, when more information about returns will be available, and W_0 is the maximum value of the investment opportunity in the first period not accounting for the option.

The investment in this simple example could be in a new fossil fuel plant for generating electricity, thus increasing greenhouse gas emissions, or conversely in reducing emissions by building a plant to generate electricity from renewable sources. Once it is recognized that an investment can be deferred until better information about the returns, including the environmental damages, or the reduction in damages, that would result, is available, it may pay to defer even though some benefit from having the investment earlier on is foregone.

Endnotes

[1] A thorough exposition of the theory of real and financial options, including stochastic processes and stochastic calculus, with many applications to economic problems, is given in Dixit and Pindyck, 1994).

[2] Following Epstein (1980), a considerable literature has grown up around the question of sufficient, and, perhaps, necessary, conditions for this result, or one like it, to hold. A very recent review is in Fisher, Hanemann, and Narain (2001). Most of the studies are quite technical, but a reassuring result is that, if benefits are simply summed over time periods, with no restriction on the relationship of one period's benefits to those of another, as in standard specifications of present value, the irreversibility effect holds.

References

Albers, H.J., and M.J. Goldbach. 2000. "Irreversible Ecosystem Change, Species Competition, and Shifting Cultivation." *Resource and Energy Economics* 22(3): 261–280.

Albers, H.J., A.C. Fisher, and W.M. Hanemann. 1996. "Valuation and Management of Tropical Forests: Implications of Uncertainty and Irreversibility." *Environmental and Resource Economics* 8: 39–61.

Arrow, K.J., and A.C. Fisher. 1974. "Environmental Preservation, Uncertainty, and Irreversibility." *Quarterly Journal of Economics* 88(1): 312–319.

Azar, C., and T. Sterner. 1996. "Discounting and Distributional Considerations in the Context of Global Warming." *Ecological Economics* 19: 169–184.

Brennan, M.J., and E.S. Schwartz. 1985. "Evaluating Natural Resource Investments." *Journal of Business* 58: 135–157.

Conrad, J.M. 1980. "Quasi-option Value and the Expected Value of Information." *Quarterly Journal of Economics* 95: 813–820.

______. 1997. "Global Warming: When to Bite the Bullet." *Land Economics* 73: 164–173.

Dasgupta, P. 2000. "Discounting: Why and How?" Tenth Annual Conference of the European Association of Environmental and Resource Economists, Rethymonon, Crete (June 30–July 2).

Dixit, A.K., and R.S. Pindyck. *Investment Under Uncertainty* (Princeton, NJ: Princeton University Press, 1994).

Epstein, L.G. 1980. "Decision Making and the Temporal Resolution of Uncertainty." *International Economic Review* 21: 269–283.

Farzin, Y., and O. Tahvonen. 1996. "Global Carbon Cycle and the Optimal Time Path of a Carbon Tax." *Oxford Economic Papers* 48: 515–536.

Fisher, A.C., and W.M. Hanemann. 1986. "Environmental Damages and Option Values." *Natural Resource Modeling* 1: 111–124.

Fisher, A.C., and U. Narain. 2001. "Global Warming, Endogenous Risk, and Irreversibility." Working Paper No. 908, Department of Agricultural and Resource Economics, University of California at Berkeley.

Fisher, A.C., J.V. Krutilla, and C.J. Cicchetti. 1972. “The Economics of Environmental Preservation.” *American Economic Review* LXII (4).

Hanemann, W.M. 1989. “Information and the Concept of Option Value.” *Journal of Environmental Economics and Management* 16: 23–37.

Henry, C. 1974. “Investment Decisions Under Uncertainty: The Irreversibility Effect.” *American Economic Review* 64: 1066-1012.

Hoel, M. 1997. “Coordination of Environmental Policy for Transboundary Environmental Problems.” *Journal of Public Economics* 66: 199–224.

Joos, F., G. Muller-Furstenberger, and G. Stephan. 1999. “Correcting the Carbon Cycle Representation: How Important is it for the Economics of Climate Change?” *Environmental Modeling and Assessment* 4: 133–140.

Kerr, R.A. 1998. “West Antarctica’s Weak Underbelly Giving Way?” *Science* 281: 499–500.

Kolstad, C.D. 1996a. “Fundamental Irreversibilities in Stock Externalities.” *Journal of Public Economics* 60: 221–233.

______. 1996b. “Learning and Stock Effects in Environmental Regulations: The Case of Greenhouse Gas Emissions.” *Journal of Environmental Economics and Management* 31: 1–18.

Krutilla, J.V., and A.C. Fisher. *The Economics of Natural Environments: Studies in the Valuation of Commodity and Amenity Resources* (Washington, D.C.: Resources for the Future, 1975), 292.

Maier-Reimer, E., and K. Hasselman. 1987. “Transport and Storage of CO_2 in the Ocean: An Inorganic Ocean-Circulation Carbon Cycle Model.” *Climate Dynamics* 2: 63–90.

Manne, A.S., and R.G. Richels. 1999. “The Kyoto Protocol: A Cost-Effective Strategy for Meting Environmental Objectives?” In “The Costs of the Kyoto Protocol: A Multi-Model Evaluation,” J.P. Weyant, editor, Special Issue of *The Energy Journal*.

Nordhaus, W.D. 1993. “Rolling the ‘DICE’: An Optimal Transition Path for Controlling Greenhouse Gases.” *Resource and Energy Economics* 15(1): 27–50.

______. 1994. *Managing the Global Commons* (Cambridge, MA: The MIT Press, 1994), 213.

Nordhaus, W.D. and J. Boyer, 2000. *Warming the World: Economic Models of Global Warming*, MIT Press, Cambridge, MA.

Peck, S.C., and T.J. Teisberg. 1992. “Ceta: A Model for Carbon Emissions Trajectory Assessment.” *Energy Journal* 13: 55–77.

Ramsey, F.P. 1928. “A Mathematical Theory of Saving.” *The Economic Journal*: 543–559.

Shultz, P.A., and J.F. Kasting. 1997. “Optimal Reductions in CO_2 Emissions.” *Energy Policy* 25: 491-500.

Ulph, A., and D., Ulph. 1997. “Global Warming, Irreversibility and Learning.” *The Economic Journal* 107: 636–650.

Weitman, M.L. 1998. “Why the Far-Distant Future Should be Discounted at Its Lowest Possible Rate.” *Journal of Environmental Economics and Management* 36(3): 201–208.

______. 2001. “Gamma Discounting.” *The American Economic Review* 91(1): 260–271.