MIT Joint Program on the Science and Policy of Global Change

The Curious Role of "Learning" in Climate Policy: Should We Wait for More Data?

Mort D. Webster

Report No. 67
October 2000

The MIT Joint Program on the Science and Policy of Global Change is an organization for research, independent policy analysis, and public education in global environmental change. It seeks to provide leadership in understanding scientific, economic, and ecological aspects of this difficult issue, and combining them into policy assessments that serve the needs of ongoing national and international discussions. To this end, the Program brings together an interdisciplinary group from two established research centers at MIT: the Center for Global Change Science (CGCS) and the Center for Energy and Environmental Policy Research (CEEPR). These two centers bridge many key areas of the needed intellectual work, and additional essential areas are covered by other MIT departments, by collaboration with the Ecosystems Center of the Marine Biology Laboratory (MBL) at Woods Hole, and by short- and long-term visitors to the Program. The Program involves sponsorship and active participation by industry, government, and non-profit organizations.

To inform processes of policy development and implementation, climate change research needs to focus on improving the prediction of those variables that are most relevant to economic, social, and environmental effects. In turn, the greenhouse gas and atmospheric aerosol assumptions underlying climate analysis need to be related to the economic, technological, and political forces that drive emissions, and to the results of international agreements and mitigation. Further, assessments of possible societal and ecosystem impacts, and analysis of mitigation strategies, need to be based on realistic evaluation of the uncertainties of climate science.

This report is one of a series intended to communicate research results and improve public understanding of climate issues, thereby contributing to informed debate about the climate issue, the uncertainties, and the economic and social implications of policy alternatives. Titles in the Report Series to date are listed on the inside back cover.

Henry D. Jacoby and Ronald G. Prinn,
Program Co-Directors

For more information, contact the Program office:

MIT Joint Program on the Science and Policy of Global Change

Postal Address: 77 Massachusetts Avenue
MIT E40-271
Cambridge, MA 02139-4307 (USA)

Location: One Amherst Street, Cambridge
Building E40, Room 271
Massachusetts Institute of Technology

Access: Telephone: (617) 253-7492
Fax: (617) 253-9845
E-mail: *globalchange@mit.edu*
Web site: *http://mit.edu/globalchange/*

Printed on recycled paper

The Curious Role of "Learning" in Climate Policy: Should We Wait for More Data?

Mort D. Webster[†]

Abstract

Given the large uncertainties regarding potential damages from climate change and the significant but also uncertain costs of reducing greenhouse emissions, the debate over a policy response is often framed as a choice of either acting now or waiting until the uncertainty is reduced. Implicit behind the "wait to learn" argument is the notion that the ability to learn in the future necessarily implies that less restrictive policies should be chosen in the near-term. I demonstrate in the general case that the ability to learn in the future can lead to either less restrictive or more restrictive policies today. I also show that the initial decision made under uncertainty will be affected by future learning only if the actions taken today change the marginal costs or marginal damages in the future. Without this interaction, learning has no effect on what we do today, regardless of what we learn in the future. Results from an intermediate-scale integrated model of climate and economics indicate that the choice of current emissions restrictions is independent of whether or not uncertainty is resolved before future decisions, because the cross-period interactions in the model are minimal. Indeed, most climate and economic models fail to capture potentially important cross-period interaction effects. I construct a simple example to show that with stronger interactions, the effect of learning on initial period decisions can be more important.

Contents

1. Introduction 1
2. An Analytical Model of Learning 4
    2.1 *Cost Functions with no Cross-Period Interaction* 6
    2.2 *Cross-Period Interaction* 7
    2.3 *The Direction of the Learning Effect* 9
3. Effect of Learning in Integrated Assessment Models 11
    3.1 *The MIT Integrated Global System Model* 11
    3.2 *Sequential Decision Using the IGSM* 12
    3.3 *The Influence of Learning in the IGSM* 14
    3.4 *Magnitude of the Learning Effect in the IGSM* 16
    3.5 *Example of an Inter-period Interaction: Induced Technical Change* 18
4. Implications for Policy 19
References 21

1. INTRODUCTION

International agreement on steps to mitigate greenhouse gases, to reduce the threat of global climate change, continues to be elusive. One characteristic of climate change that makes consensus difficult is the magnitude of the uncertainty regarding both the costs and impacts. The amount of climate change that may occur and the effects resulting from such a change are potentially very large, including changes in precipitation patterns, sea level rise, frequency and severity of extreme climatic events, and even a shift in the ocean currents that warm Europe.

[†] The author is a post-doctoral associate with the Massachusetts Institute of Technology, Joint Program on the Science and Policy of Global Change: 77 Massachusetts Ave., Bldg. E40-271, Cambridge, MA 02139-4307, USA.

But uncertainty about these effects remains very large. The costs of restricting greenhouse gas emissions are also uncertain, and are estimated by some to also be quite large as well.

Greenhouse gases are an example of what, in economic terms, is known as a “stock” pollutant. Once emitted into the atmosphere, these gases have extremely long lifetimes, hundreds to thousands of years for some of them. Their effects on the climate system are a function of their total concentrations in the atmosphere, which change slowly as a result of emissions over many decades. The uncertainties in climate change, the long-time scales involved, and the potentially irreversible effects—combined with the fact that control policies taken today can be reconsidered later—make the climate change issue one of deciding what to do now given that we may resolve some of the uncertainties in the future. Thus, discussions of climate policy are typically framed as a choice of either acting to reduce emissions now or waiting until we learn more about the problem. Stringent and costly actions taken now might prove to be unnecessary if the climate change problem turns out to be not as bad as we thought. On the other hand, we may regret not acting aggressively now if we learn that the effects of climate change are much more severe than expected. Researchers and interest groups alike have made both cases (*e.g.*, Risbey *et al.*, 1991a, 1991b; Schlesinger and Jiang 1991a, 1991b; Stevens, 1997; United Nations, 1992). In the policy debate the most common argument is that the expectation of future learning should lead to less action now than otherwise (the so-called “wait to learn” argument). In the available economic literature, however, there is no consensus on the issue.

One stream of research on decision-making under uncertainty within economics has focused on the additional value of avoiding future damages when those damages are uncertain and irreversible, and so concludes that the ability to learn should lead to lower emissions if those emissions has irreversible consequences (*e.g.*, Arrow and Fisher, 1974; Henry, 1974; Chichilnisky and Heal, 1993). Others have reached the opposite conclusion, that the ability to learn should lead to higher emissions, because of irreversibility in the long-lived capital stock (*e.g.*, Viscusi and Zeckhauser, 1976; Ulph and Ulph, 1997; Pindyck, 1999). The most general result from the theoretical literature to date is from Epstein (1980), who showed that learning in the presence of an irreversibility can lead to either more or less of the irreversible development activity. For Epstein, direction of the effect depends on the shape of the marginal cost function (*i.e.*, the derivative of the objective function): if the marginal cost is concave then learning leads to less of the activity, and if it is convex then learning leads to more of the activity. Unfortunately, requiring strict concavity or convexity is overly restrictive for representing climate change, as shown by Ulph and Ulph (1997). Further, Epstein’s result does not address the conditions under which learning has no effect (*i.e.*, cases with and without learning lead to the same level of activity in period 1), or what determines whether the magnitude of a divergence is large or small.

Also, a number of studies of uncertainty and decision-making in the climate issue use integrated economic-climate models. Several of these do address uncertainty, but do not consider the influence of learning on the near-term decision. Rather, they focus on related but distinct questions. Examples include the optimal decision under uncertainty when the uncertainty will later be resolved (Hammitt, Lempert, and Schlesinger, 1992), the comparison between choice under perfect certainty and choice under uncertainty that is later resolved (Manne and Richels, 1992, 1995), and the value of perfect information revealed at different times in the future (Nordhaus, 1994; Nordhaus and Popp, 1997).

Studies that have explicitly examined the effect of learning in empirical models of climate change (Nordhaus, 1994; Kolstad, 1996; Ulph and Ulph, 1997) have found that learning seems to have almost no effect on the period 1 strategy. The explanations in these studies for the lack of an effect of learning rely on two characteristics of the models:

1) The irreversibility constraint does not bind: *i.e.*, the damage losses are not severe enough to drive period 2 emissions to zero[1] (Kolstad, 1996; Ulph and Ulph, 1997) and,

2) The stock nature of greenhouse gases: the fact that the existing stock decays very slowly means that period 1 emissions have very little influence on the total stock of greenhouse gases in the atmosphere (Kolstad, 1996; Nordhaus, 1994).

Further, results from empirical models are contradictory, with learning leading to more period 1 abatement in results from Nordhaus (1994) but less abatement in results from Ulph and Ulph (1997).

In this paper, I will clarify the effect of learning by representing the process of sequential choice, with the possibility of learning, in a simplified way with a two-period decision.[2] The first period represents “now”: the decision that will be implemented over the next few years. The second period represents the notion that we can do something different “later,” whether we have reduced uncertainty or not. For this two-period decision under uncertainty, the question is: how does the optimal first period strategy change if uncertainty is resolved before the second period decision is made? The influence of learning is examined by considering two extreme cases: 1) ”No Learning,” in which the uncertainty at the time of the period 2 decision is the same as for period 1; and 2) “Complete Learning,” in which all uncertainty is resolved before the period 2 decision.

These two cases are illustrated in **Figure 1**. The upper part of the figure shows the Complete Learning case, where the decision is made today about what level of emissions constraints to implement; then after ten years the uncertainty about climate impacts is eliminated, and then the decision about increasing or relaxing emissions constraints is made with perfect knowledge. In contrast, the lower panel in Figure 1 shows the No Learning case where the decision today and the decision in ten years are both made under the same level of uncertainty about climate change impacts. The effect of learning can be seen by comparing the best decision made today in each of these two cases.

---

[1] According to Kolstad and Ulph and Ulph, the irreversibility of the period 1 decision is only binding if, when the true state of nature is revealed in period 2 through learning, one would wish to undo the action in period 1. If damages are not severe enough to warrant a decision to “negatively emit” (take carbon out of the atmosphere), then the irreversibility of emissions in period 1 becomes irrelevant.

[2] A two-period model of sequential decision gives insight into the general effect on the period 1 decision that would be obtained from a model with three, four, or more decision points.

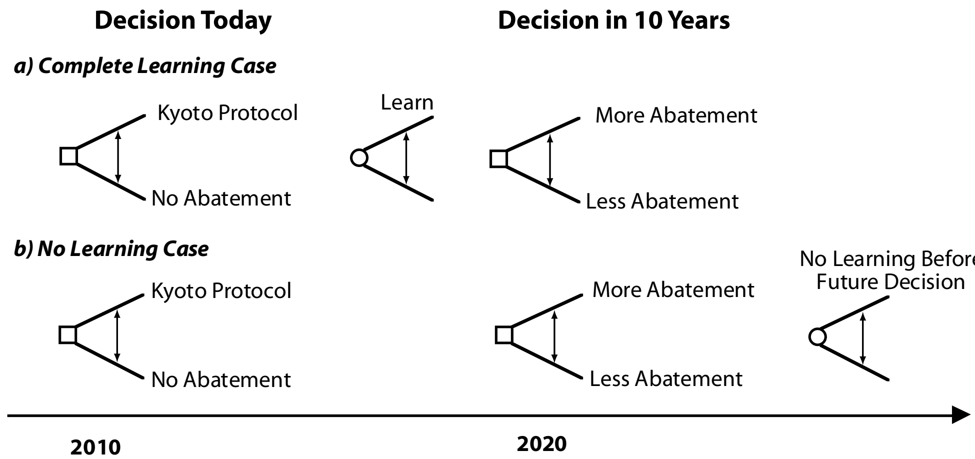

Figure 1. Policy Choice as Two-period Decision with and without Learning

The analysis below demonstrates the conditions for the existence of an effect by learning on strategy choice, and explains the factors that determine the direction of the learning effect. In Section 2, I develop general results using an analytical dynamic programming model of a two-period decision. Then, in Section 3, I demonstrate the effect of learning with an integrated assessment model that represents much of the complexity in the economic and climate systems. Section 4 summarizes the main findings.

2. AN ANALYTICAL MODEL OF LEARNING

The choice of a climate change policy under uncertainty can be defined as a dynamic programming problem, with 2 periods. Define

$t = 1, 2$

$X_t \equiv$ the set of all possible emissions levels that can be chosen in period t

$x_t \equiv$ the level of emissions allowed in period t, chosen from the set X_t

$\theta \equiv$ the severity of damage costs from climate change

$C_1(x_1, \theta) \equiv$ abatement costs and damage costs in period 1

$C_2(x_1, x_2, \theta) \equiv$ abatement costs and damage costs in period 2

and let $E_\theta\{\cdot\}$ denote the expectation with respect to the marginal distribution of θ.

In each period t, a level of allowed emissions x_t is chosen. Each period has a total cost function $C_t(\cdot)$ stated in terms of emissions level, which includes both abatement costs and damages from the accumulated stock of carbon. The uncertainty in the damages from climate change is represented by different states of the world θ that may obtain, and so the damage costs (and therefore the total costs) are also a function of θ.

Next, define two value functions, one representing the minimized costs over both periods, and the other the minimized costs in period 2 only:

$V_1(x_1, x_2, \theta) \equiv$ the sum of abatement costs and damage costs over both periods given emissions level x_1 at t = 1, emissions level x_2 at t = 2, $\tilde{\theta} = \theta$,

$V_2(x_1, x_2, \theta) \equiv$ the abatement and damage costs at t = 2 if emissions level x_2 is chosen in light of $\tilde{\theta} = \theta$, given that the emissions level at t = 1 was x_1.

To evaluate the effects of learning, consider two extreme cases. In the "learning" case, the uncertainty in the damage costs θ is completely resolved before the period 2 decision is made. In the "no learning" case, the uncertainty in θ will be the same in both periods. By comparing the optimal period 1 strategy in these two cases, we will see the maximum possible effect of learning on strategy. For this study, we consider learning to be autonomous or exogenous, in which the true state is revealed with the passage of time. Other approaches to modeling learning include active learning, in which the evolution of the climate and economy are observed and beliefs are updated, and purchased learning, in which improved information is purchased with an explicit cost that is modeled (*e.g.*, R&D). Exploration of those cases is a logical extension of this work. In this model, learning only resolves the uncertainty in damage costs, while abatement costs are treated as certain.

A dynamic programming problem is always solved through backward induction—*i.e.*, by finding the optimal choice for the last decision first, and then working backwards. For this problem, we first solve for the optimal decision in period 2 as a function of the period 1 strategy. Then we substitute this expression into the objective function and solve for the optimal first period strategy. The first step for the problem posed here is to choose, at t = 2, the emissions level x_2^* that minimizes the sum of abatement and damage costs given that period 1 emissions level was x_1 and that the severity of climate change damages is θ. In the "learning" case, the optimal emissions level x_2^* is chosen with certainty about θ, and the value function for period 2 is

$$V_2(x_1, x_2, \theta) = \min_{x_2 \in X_2} [C_2(x_1, x_2, \theta)]. \quad (1a)$$

In the "no learning" case, the optimal emissions level x_2^* must be chosen under uncertainty in θ, and the value function for period 2 is

$$V_2(x_1, x_2, \theta) = \min_{x_2 \in X_2} [E_\theta\{C_2(x_1, x_2, \theta)\}]. \quad (1b)$$

Without learning, the best strategy is the one that minimizes the expected value of the costs.

Once the second period optimal strategy, x_2^*, is found, the next step is to substitute this expression into the value function V_1, and solve for the optimal period 1 strategy. We must choose the optimal period 1 emissions level x_1^* that minimizes the expectation of the sum of costs over both periods. Namely,

$$V_1(x_1, \tilde{\theta}) = \min_{x_1 \in X_1} [E_\theta\{C_1(x_1, \tilde{\theta}) + V_2(x_1, x_2^*(x_1, \tilde{\theta}), \tilde{\theta})\}] \ . \quad (2)$$

There are several important characteristics of this abstract model worth highlighting. First, the stock nature of the problem is represented by the dependency of C_2, the cost in the second period, on x_1, the decision made in the first period. In multi-period decisions about stock pollutants, capital stock, or other quantities that accumulate over time, the costs and/or benefits in any period are partly a function of decisions made in previous periods. Analogously, the current period's decisions will have cost/benefit impacts in future periods. This formulation is in contrast to flow-type problems in which the implications of each period's decision are felt in that period only, and costs have no relation to what has previously occurred (*e.g.*, noise pollution).

The second element of the model to note is the difference between the learning case and the no learning case. In the no learning case, there is only a single choice of x_2^* that must minimize the mean or expected costs across all possible states of the world, since we don't yet know which one is the true state. In contrast, in the learning case many different optimal choices of x_2^* exist, each minimizing costs in the particular state of the world θ_i. Of course, even when there is learning by period 2, the period 1 decision must still be made based on the expected value over all possible states, as can be seen by equation 2.

The effect of learning can be evaluated as the difference between x_1^L, the solution to Eq. 2 when uncertainty is completely resolved (Eq. 1a), and x_1^N, the solution to Eq. 2 when no learning occurs (Eq. 1b). With the generic form of this model, we cannot say much about what these two expressions might be, and so in the following sections we consider alternative functional forms of the cost functions.

2.1 Cost Functions with no Cross-Period Interaction

Begin by assuming that the period 2 cost function has no cross-products between the period 1 and period 2 strategies. In constructing the cost functions we make the simple assumptions that the first period costs, given a state of the world θ, are linear,

$$C_1(x_1,\theta) = a(\theta)x_1 + b(\theta)\,, \tag{3}$$

and that the second period costs are a simple quadratic function of both periods' decisions,

$$C_2(x_1,x_2,\theta) = c(\theta)x_2^2 + d(\theta)x_2 + e(\theta)x_1^2 + f(\theta)x_1 + g(\theta)\,. \tag{4}$$

In this model, the stock nature of the problem is represented by the terms $e(\theta)x_1^2$ and $f(\theta)x_1$ in the second period cost function. The decision made in the first period will influence costs in the second period.[3]

The first step is to solve the second period decision. There are two cases: one with learning and one without. It can be shown (Webster, 2000) that the optimal period 2 strategy for the case where uncertainty is resolved is:

$$x_2^{*L} = -\frac{d(\theta)}{2c(\theta)}\,. \tag{5}$$

Because the period 2 emissions level will be chosen after we know the true state of the world (*i.e.*, we will know the values of the coefficients *a* through *g* of Eqs. 3 and 4), the optimal strategy is a

[3] These terms do not capture any change in the marginal damages that may occur with a non-linear damage function. In this formulation the total period 2 damage is a function of period 1 strategy, but the marginal damage is not. The dependence of marginal damage on first period strategy is a different effect, and is treated in the next example in Section 3.2.

function of those values. In contrast, the optimal period 2 strategy without learning will be a function of the expectation of these coefficients, since their true values will still be uncertain:

$$x_2^{*N} = -\frac{\bar{d}}{2\bar{c}}. \tag{6}$$

Having solved for the optimal period 2 strategy in each case, we substitute for x_2 in Eq. 4, and then solve for the optimal emissions level x_1^* that minimizes costs over both periods. It also can be shown (Webster, 2000) that the optimal period 1 emissions when learning occurs is

$$x_1^{*L} = -\frac{\bar{a}+\bar{f}}{2\bar{e}} \tag{7}$$

and that the optimal strategy in period 1 without learning is also

$$x_1^{*N} = -\frac{\bar{a}+\bar{f}}{2\bar{e}}. \tag{8}$$

The two solutions for this example are identical: $x_1^L = x_1^N$. Although the period 1 decision does affect the *total* costs in period 2 (damages from the remaining stock), it does not interact in any way with the period 2 decision through an influence on *marginal* cost. The lack of dependency of period 2 decisions on period 1 actions is clear from Eqs. 5 and 6, since x_1 does not appear in either solution. This result leads to a more general proposition. For any two-period sequential decision under uncertainty represented by Eqs. 1 and 2, and where the solution to Eq. 1a is denoted by x_1^N and the solution to Eq. 1b is denoted by x_1^L,

$$\textit{If} \ \frac{\partial^2 C_2}{\partial x_1 \partial x_2} = 0 \ \textit{then} \ x_1^L = x_1^N.$$

The proof of this proposition is given in Webster (2000).

The basic conclusion is that if today's decision has no effect on the marginal costs of tomorrow's decision, then the two choices are independent. Whether we learn or not (which does influence tomorrow's decision) is irrelevant for today's decision. Today's decision is merely made on the basis of costs and benefits today, plus the expected discounted costs and benefits that continue to accrue in future periods as a result of today's decision (the non-interacting stock effects).

2.2 Cross-Period Interaction

The previous example showed that if the cost function for period 2 has no interaction terms between the strategies in the two periods, then the period 1 optimal choice is unaffected by assumptions about learning. What happens when there is an interaction term? To construct this case, assume the same linear cost function for period 1 as in the previous example,

$$C_1(x_1,\theta) = a(\theta)x_1 + b(\theta). \tag{9}$$

Define the period 2 cost function with a quadratic term in the second period strategy and a single linear interaction term between the two decisions:

$$C_2(x_1,x_2,\theta) = c(\theta)x_2^2 + d(\theta)x_1x_2 \tag{10}$$

This functional form offers the simplest formulation with a cross-period effect, dropping other terms that appeared in Eq. 4. In particular, the linear and quadratic terms in x_1 that represent the non-interacting stock effect from period 1 decisions are omitted. In the previous example these terms, which influence total but not marginal costs, were shown to have no effect on the period 1 decision.

To consider the effect of learning, we begin, as above, by solving for the optimal period 2 strategy. For the learning case, the cost-minimizing strategy is:

$$x_2^{*L}(\theta) = -\frac{d(\theta)}{2c(\theta)} x_1 . \tag{11}$$

The optimal strategy when no learning occurs is:

$$x_2^{*N} = -\frac{\bar{d}}{2\bar{c}} x_1 . \tag{12}$$

In contrast to the previous example, the optimal strategy in period 2 here depends on the period 1 strategy (x_1) that was chosen and is directly due to the $x_1 x_2$ cross-product term in Eq. 10.

Substituting these expressions for x_2 in Eq. 10, the cost-minimizing strategy in period 1 when learning will occur is:

$$x_1^{*L} = \frac{2\bar{a}}{E_\theta \left\{ \dfrac{d^2(\theta)}{c(\theta)} \right\}} . \tag{13}$$

The cost-minimizing period 1 emissions level in the no learning case is:

$$x_1^{*N} = \frac{2\bar{a}\bar{c}}{\bar{d}^2} . \tag{14}$$

Adding the term $d(\theta)x_1x_2$ to Eq. 10 has caused period 1 strategy to now depend on learning. The solutions have a common component, $2\bar{a}$, scaled by the expectation of a non-linear function, d^2/c, in the case of learning, and by the non-linear function of expectations in the case of no learning.

What might the cross-period interaction[4] term $d(\theta)x_1x_2$ in Eq. 10 represent? As noted earlier, it could represent non-linearity in the damage function. If a larger stock of CO_2, resulting from higher period 1 emissions, changes the marginal damages in period 2, then this change in marginal damages will show up in this cross-term. It could also represent a dependency of the marginal cost of mitigation in period 2 on period 1 decisions (*e.g.*, capital stock effects). Depending on the sign, the coefficient of the cross-term represents the fact that decisions in different time periods can act as substitutes ($d>0$) or as complements ($d<0$). Decisions act as substitutes when an *increase* in the period 1 activity level, x_1, results in an additional *increase* in the marginal cost of choosing higher level of x_2 in period 2. Generally substitution exists in problems where there is some finite resource that can be used across both periods; using more of the resource in period 1 leaves less of the resource or results in increasingly expensive

[4] The "interaction" described here, the dependence of marginal costs in period 2 on the period 1 strategy, is independent of learning. This interaction is present even in the no learning case. Thus it is a different phenomenon than "interactive learning" in the tradition of models of learning-by-doing (*e.g.*, Miller and Lad, 1984).

alternatives for the second period decision. Decisions act as complements when an *increase* in activity in the first period causes the per-unit cost of period 2 action to *decrease*. Complementary situations typically exist when first period action constitutes some form of investment that reduces future costs.

A special case of the solutions (Eqs.) 13 and 14 will be useful in Section 3. Consider the discrete distribution case where there are two possible states:

$$\theta = \{low, high\}.$$

Denote the cost coefficients as

$$a(low) = a_L; \qquad a(high) = a_H \ .$$

We also define

$$P \equiv \Pr\{\theta = high\},$$

the probability of being in the *high* damage state of the world. Thus the expectation with respect to θ can be written:

$$\overline{a} = E_\theta \{a(\theta)\} = Pa_H + (1-P)a_L \ .$$

For this simple case with only two discrete states, it can be shown that the optimal period 1 strategies become

$$x_1^N = 2\overline{a}\left(\frac{\overline{c}}{\overline{d}^2}\right), \tag{15}$$

$$x_1^L = 2\overline{a}\left(\frac{c_L c_H}{c_L d_H^2 P + c_H d_L^2 (1-P)}\right), \tag{16}$$

for the cases without and with learning, respectively. Note that only the bracketed term differs between Eqs. 15 and 16. For this problem, the optimal period 1 decision is determined by the average or expected period 1 marginal cost $2\overline{a}$, scaled by a term representing the second period marginal costs. In Section 3 we will use these expressions to demonstrate the magnitude of the divergence between strategies exhibited by an empirical climate assessment model.

2.3 The Direction of the Learning Effect

The previous examples show that a cross-period interaction is a necessary condition for learning to influence period 1 strategy. However, the presence of an interaction is not a sufficient condition. Further, if learning does influence the period 1 choice, does it result in higher or lower emissions? Even without deriving the expressions that demonstrate the determinants of the direction of the learning effect (see Webster, 2000), it is possible to provide some intuition on how learning might either increase or decrease the optimal emissions level in period 1.

In the simple model with aggregate total (damage plus abatement) cost functions defined by Eqs. 9 and 10, learning always leads to a lower strategy level x_1^* because costs change monotonically with strategy level. There is a downside to doing too much, but no equivalent downside to doing too little[5]. Thus the irreversibility and uncertainty leads to a lower level of the

[5] Or the opposite, depending on the signs of the coefficients.

irreversible activity if learning and correction are possible later. If we consider the costs as damage losses and x_t as emissions, this is equivalent to the Arrow and Fisher (1974) result that learning leads to lower emissions in period 1. If we consider the costs as representing only abatement costs and x_t as emissions level, this is equivalent to the investment under uncertainty models of Pindyck (1991) in which learning will lead to higher emissions (*i.e.*, less abatement with learning).

Real-world problems such as a decision about climate policy involve both abatement and damage costs that change in opposite directions with the emissions level. A more general representation is to treat abatement costs (decreasing in emissions) and damage costs (increasing in emissions) separately. For a two-stage decision in which learning resolves uncertainty in damage costs, *the optimal emissions strategy chosen with learning may be higher or lower than without learning*. The irreversibility in both damages and control costs causes two effects from learning, pulling in opposing directions. Learning, thereby, can lead to higher or lower emissions depending on the relative magnitudes of the control costs and damage costs.

The dominant direction of the learning bias can be explained in terms of two elements of the decision: 1) the anticipation of the period 2 strategy, and 2) the regret over the period 1 choice given the outcome after learning. When a period 1 strategy is chosen under uncertainty, and then the uncertainty is resolved in period 2, some regret over the period 1 choice is inevitable[6]. Suppose in period 2 we learn that the damages from climate change are less severe than expected. Then we will choose a higher level of emissions in period 2 than we did under no learning. Because of the interaction, we will wish that we had anticipated this higher emissions level in period 2, and also chosen higher emissions in period 1. We will regret having spent more on abatement cost in period 1 than turned out to be necessary. Now, suppose instead we learn in period 2 that the damages are more severe than the expectation. In this case, we will lower emissions further in period 2 than we would have without learning. Because of the interaction, we will wish we had anticipated this lower emissions strategy by emitting less in period 1. We will regret not having taken enough precaution in the face of uncertain climate damage.

The net effect of learning on strategy is determined by the relative magnitudes of these two regrets. When the decision is being made in period 1, we are still uncertain whether we will learn that damages are greater or less than the expectation. The probability distribution over damage costs reflects our belief in the relative likelihood of each state of the world that might be revealed in period 2. If, on balance, the dominant regret will be that we will have spent too much on abatement before learning that damage costs are lower, then learning will lead to a net increase in period 1 emissions. We call this the "sunk cost" situation. If, on the other hand, the dominant regret over all possible outcomes will be that we should have abated more, then learning will lead to a net decrease in emissions in period 1. We call this situation the "precautionary case." If the regret from abating emissions when damages are revealed to be low and the regret from

[6] Except, of course, in the rare case that the revealed true state is exactly equal to the expectation under uncertainty.

abating too little when damages are revealed to be high balance each other, then learning may still not influence the period 1 decision, even in the presence of an interaction.

Thus, although the direction of learning effect is influenced by the convexity or concavity of the marginal costs (Epstein, 1980), it also is determined by the shape of the probability distribution over uncertain abatement costs and damage costs. When the expected damages from climate change are low (*i.e.*, low expected net benefits) but there is a small probability of high damage cost (*i.e.*, skewed towards high damages), learning will lead to lower period 1 emissions. The regret from learning that damages are high will dominate, since the low damages were already the expectation. Conversely, if expected damage costs are high, but there is a small probability of low damage costs (*i.e.*, skewed towards low damages), then learning will lead to higher emissions in period 1. We will demonstrate the dependence on the probability distribution of the direction of the learning effect with results from the integrated assessment model in the next section.

3. EFFECT OF LEARNING IN INTEGRATED ASSESSMENT MODELS

3.1 The MIT Integrated Global System Model

The analytical model used above to explore the effect of learning employed highly simplified cost functions. In this section the effect of learning is illustrated using a climate policy assessment model of intermediate complexity. The integrated assessment model used is the MIT Integrated Global System Model [IGSM] (Prinn *et al.*, 1999), augmented with a damage function related to change in global mean temperature. The economic component of the model, the Emissions Projections and Policy Analysis (EPPA) model (Babiker, 2000) is a recursive-dynamic computable general equilibrium model, consisting (in the calculation applied here) of twelve geopolitical regions linked by international trade, ten production sectors in each region, and four consumption sectors. The climate component is a two-dimensional (zonal averaged) representation of the atmosphere and (Sokolov and Stone, 1998). The climate model includes parameterizations of all the main physical atmospheric processes, and is capable of reproducing many of the non-linear interactions simulated by atmospheric GCMs.

In order to choose one set of strategies as "optimal," we require a basis for comparing the costs of reducing emissions with the benefits of avoiding damages. We augment the EPPA mitigation cost model with the Nordhaus damage function (Nordhaus, 1994). This damage function has been widely used (*e.g.*, Kolstad, 1996; Lempert *et al.*, 1996; Peck and Teisberg, 1992; Pizer, 1999), and facilitates the comparison of results here with other studies. The Nordhaus damage function estimates the percentage loss of gross world product as a function of the global mean temperature change,

$$d(t) = \eta[\Delta T(t)]^{\pi} \tag{17}$$

where $d(t)$ is the fraction of world product lost due to climate damages in year t, and $\Delta T(t)$ is the increase in global mean temperature from preindustrial levels.

Solving for an optimal sequential decision under uncertainty requires a large number of simulations of the empirical economic-climate model. However, the IGSM requires too much computation time for this many simulations to be feasible. To perform the calculations, we estimated reduced-form versions of the IGSM using the Deterministic Equivalent Modeling Method (Tatang *et al.*, 1997; Webster and Sokolov, 2000). These simpler functional forms have been shown to replicate the results of the original IGSM to within a 1% error of the mean (Webster, 2000). The reduced-form models are used in all calculations below.

To set up the sample calculation, the sequential decision problem for climate change is defined in Section 3.2. The effect of learning in the IGSM is described in Section 3.3, and then in Section 3.4, results from the analytical model are used to explain the effect of learning in this model. Finally, in Section 3.5, I use induced technical change as an example of a strong interaction which can be added to the model to increase the effect of learning.

3.2 Sequential Decision Using the IGSM

As above, we frame the control choice as a sequential decision under uncertainty, with the two decision points as illustrated in Figure 1. The decision-maker in this model represents the aggregate "Annex I," the industrialized nations that would constrain emissions under the Kyoto Protocol. The objective function for the decision-maker is to minimize the net present value of total consumptions losses. Consumption losses occur both as a result of constraining carbon emissions and as a result of the impacts of climate change. The stream of costs over time is discounted at a reference rate of 3%, which is subjected to sensitivity analysis. The set of possible strategies in this model represents choice over levels of emissions abatement only; other possible complementary policies of research, adaptation, and geoengineering are not modeled here. We assume that only Annex I nations constrain emissions in this model, while the less developed nations increase their emissions of greenhouse gases unrestricted over the entire 100-year time horizon.

The strategies are defined as maximum allowable growth rates in emissions. The first period strategy can be any rate between 0% per year (emissions stabilization) and 1.4% per year (unconstrained for all regions) over the years 2010-2019 (**Table 1**). The second period strategy is chosen from a low of –0.8% per year and a high of 1.2% per year (unconstrained), and constrains emissions for the years 2020-2100. The period 1 strategy also determines the absolute emissions level in 2010, as indicated in **Table 2**, which shows the reduction in relation to the Annex I Kyoto target (United Nations, 1997).

Table 1. *Strategy Choices in each Period: Maximum Allowable Emissions Growth*

Decision Period	*Strategy Variable*	*Years*	*Most Stringent Constraint*	*Least Stringent Constraint (No Limits on Emissions Growth)*
1	*Policy2010*	2010-2019	0% per Year	1.4% per Year
2	*Policy2020*	2020-2100	–0.8% per Year	1.2% per Year

Table 2. *Emission Targets for 2010 as a Function of Strategy Level*

Policy2010	*2010 Emissions Constraint*
0%	100% of Kyoto
0.2%	85% of Kyoto
0.4%	70% of Kyoto
0.6%	55% of Kyoto
0.8%	40% of Kyoto
1.0%	25% of Kyoto
1.2%	10% of Kyoto
1.4%	Reference (No Controls)

Based on previous work (Webster, 2000; Webster and Sokolov, 2000), we consider three uncertain parameters that have the greatest impact on damage costs:

- Climate Sensitivity: this parameter determines the change in global mean temperature at equilibrium that results from a doubling of CO_2. Different sensitivities are obtained in the MIT 2D climate model by adding an additional cloud feedback (Sokolov and Stone, 1998).
- Rate of Ocean Uptake: the 2D climate model parameterizes the mixing of both heat and carbon from the mixed-layer ocean into the deep ocean. A slower ocean will result in both higher carbon concentrations in the atmosphere and in more rapid warming (Sokolov and Stone, 1998).
- Damage Valuation: to reflect the large uncertainty in the valuation of climate change impacts, we treat the damage coefficient η from Eq. 17 as uncertain. This coefficient reflects the percentage of gross world output that would be lost from a 3˚C temperature rise (Nordhaus, 1994).

The probability distributions for the three uncertain parameters are discrete two-point approximations based on continuous distributions, and are subjected to extensive sensitivity testing. The reference continuous distributions are obtained from expert elicitation. The distributions for climate sensitivity and for ocean uptake are given in Webster and Sokolov (2000), based in part on Morgan and Keith (1995). The distribution for the damage valuation is taken from Roughgarden and Schneider (1999), based on the assessment by Nordhaus (1994b). Because the distributions are based on expert elicitation, they are subject to all the biases of subjective judgment about probability (Morgan and Henrion, 1990). Also, for almost all parameters, there is wide disagreement between experts. It is crucial therefore to subject all results from decision models to sensitivity testing of the assumed distributions. We approximate the continuous distributions with the discrete distributions shown in **Table 3**. Sensitivity testing is then performed by varying the probability of the high damage state (Branch 2).

Table 3: *Distributions for Uncertain Quantities*

	Branch 1 (P=0.8)	**Branch 2 (P=0.2)**
Climate Sensitivity (°C)	2.5	4.5
Oceanic Uptake (cm^2/s)	2.5	0.5
Damage Cost Coefficient (%)	.02	.16

As in the analytical model of Section 2, learning is modeled as the revelation of the true state of damage costs in period 2. This is the simplest representation of learning that allows the examination of the effect of learning on strategy. More sophisticated models of learning such as including explicit costs of reducing uncertainty or Bayesian updating of probability distributions from observations are left for future studies. Also, as above, learning does not resolve the uncertainty in abatement costs.

3.3 The Influence of Learning in the IGSM

Using the IGSM and the two-period sequential framing of emissions control choice, the effect of learning can be explained in this more complex setting. **Figure 2** shows the emissions for both the No Learning and the Learning cases, using the reference probability distributions given in Table 3. We find, not surprisingly, that with learning (dashed lines) the choice among the eight discrete strategies in period 2 differs from the choice in the no learning case (solid line), depending on which of the eight possible states is revealed about the three uncertain parameters. As viewed from 2000, the probability (P) of a strategy being chosen in the learning case is indicated in Figure 2. With probability P = 0.8, we will learn that the damages are not as serious as expected, and will choose higher emissions. With P = 0.14, we learn that damages are worse than expected and emit less. There is a 0.06 probability that we will choose the same emissions as in the no learning case. The main focus of this paper, however, is what happens to period 1 emissions. Here there is no difference between the learning and no learning cases, and, in fact, the optimal strategy in period 1 is to do nothing in all cases.

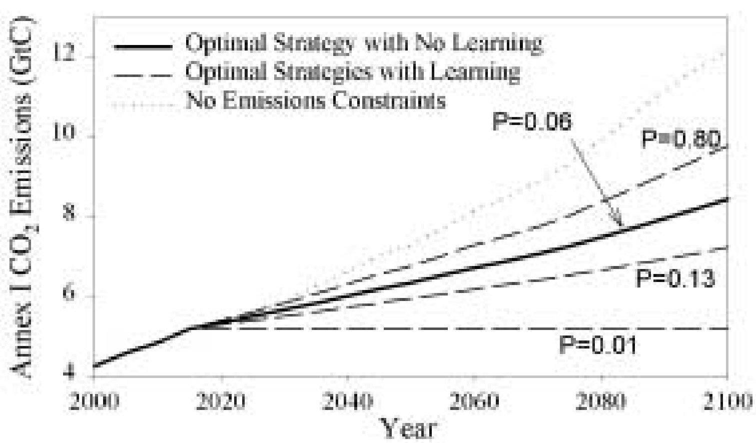

Figure 2. Optimal Annex I Emissions With and Without Learning

The result that learning has no effect on first period strategy is surprising but, as noted earlier, one consistent with other studies. Is this result dependent on the specific probability distributions used? Sensitivity analysis not reported here indicates that, of the three uncertain parameters treated, the minimized total expected losses are most sensitive to the damage valuation parameter. If the probability of high damage cost is varied from 0 to 1.0, while the probabilities of high climate sensitivity and slow ocean uptake are kept at the reference values of 0.2, the same strategy (no controls) is optimal with and without learning. For any distribution of the damage cost parameter, however, the optimal strategy in period 1 is independent of whether learning occurs.

In order to find a divergence between the optimal strategy with learning and the optimal strategy without learning, the probabilities of all three uncertain parameters must simultaneously be adjusted far from the reference values. **Figure 3** shows the optimal strategies with and without learning when the probability of high climate sensitivity (4.5˚C) is assumed to be 0.85, and the probability of slow ocean uptake (0.2 cm^2/s) is also assumed to be 0.85. Here, there is difference between the strategies chosen with and without learning when the probability of high damage cost is between 0.95 and 0.975. In this region, it is optimal to constrain emissions more if learning will occur. As the probabilities of all three uncertain parameters are varied simultaneously, other sets of assumptions will yield differences between strategy with and without learning, but only for a small fraction of possible assumptions. In addition, such regions of divergence in strategies only exist for assumptions about the uncertain parameter distributions that are inconsistent with expert judgment, as in Figure 3.

In addition to varying the probability distributions, other assumptions of the decision model can be altered in an attempt to get a stronger effect on strategy by learning. Other cases that have been tested include

- Lengthening the first period from ten to forty years (*i.e.*, first period decision determines emissions from 2010-2050),
- Assuming very slow ocean uptake with certainty (to increase the "irreversibility" of emissions),
- Varying discount rates between 0% and 10%,

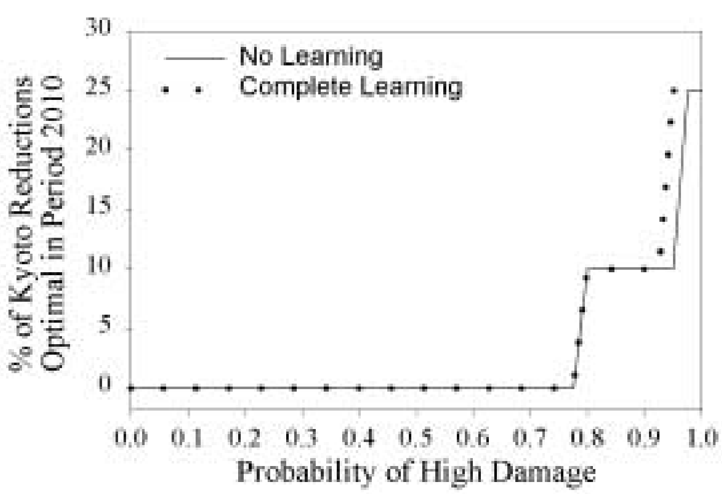

Figure 3. Optimal First Period Strategy with and without Learning. Probability of high climate sensitivity is 0.85 and probability of slow ocean uptake is 0.85.

and several other tests (Webster, 2000). In all of these cases, the optimal first period decision with learning and the optimal decision without learning are almost always the same. The ranges of uncertain parameters that result in a divergence are small, and usually occur under parameter distributions that are inconsistent with expert opinion.

3.4 Magnitude of the Learning Effect in the IGSM

The analytical model in Section 2 leads to a suspicion that the reason for this lack of influence of learning is that the inter-period interactions in this model are insignificant. And indeed, despite the complexity captured by the economic and climate models in the IGSM, few of the possible interactions over time of emissions control levels are represented in this model. The magnitude of the cross-period interactions can be estimated by examining the reduced-form models that are fit to the full IGSM. We estimate the control costs and the damage costs in each period as a function of the strategy level chosen in each of the two periods. To facilitate comparison with the analytical results from Section 2, we use the same quadratic functional form as the cost functions there, including a single linear cross-term.

Two estimated cost-functions were fit: one for low climate damage costs and one for high damage costs. The reduced-form estimates of total costs in period 2 as a function of the strategies are given in Eqs. 18 and 19 for the high damage and low damage assumptions, respectively.

$$TC_2(x_1, x_2, H) = 27108 - 2482x_1 + 153{x_1}^2 - 2648x_2 + 226{x_2}^2 + 177x_1x_2 \tag{18}$$

$$TC_2(x_1, x_2, L) = 18712 - 2601x_1 + 145{x_1}^2 - 2962x_2 + 196{x_2}^2 + 161x_1x_2 \tag{19}$$

From Eqs. 15 and 16, we know that the optimal period 1 strategies with and without learning are each equal to a common term scaled by a term that varies depending on whether learning will occur. Using the coefficients from Eqs. 18 and 19, and assuming that the probability of high damage costs $P = 0.5$ to maximize the uncertainty, the optimal period 1 strategies are proportional to:

$$x_1^L \propto \frac{c_L c_H}{c_L d_H^2 P + c_H d_L^2 (1-P)} = 0.00738 \tag{20}$$

$$x_1^N \propto \frac{E\{c\}}{E\{d\}^2} = 0.00739 \tag{21}$$

Because the ratio differs only slightly from 1.0, we expect little influence on the period 1 strategy from learning. This summary calculation is indicative of why, even across a wide range of distributions and similar tests, we find very few instances when learning affects the period 1 decision. Examining the coefficients in Eqs. 18 and 19, we see that the cross-term coefficients are not zero, but relative to the independent effects of the strategies in each period, the interactions are very weak.

What are the interactions in the IGSM represented by these coefficients? One process in the IGSM that will cause the first period strategy to affect the second period marginal costs is the vintaging of capital in the economic model component. In the EPPA model, a portion of the preexisting capital stock in any period is not malleable (cannot be shifted to different sectors when relative prices shift) and the proportions of input factors are also frozen at the current technology levels. As a result, a lack of abatement in the first period can result in investment in new carbon-emitting capital, some of which cannot be shifted if abatement is undertaken in the second period (Jacoby and Sue Wing, 1999). As a result, a less stringent policy in period 1 (higher emissions) will cause a higher marginal cost of emissions reductions in the second period. Another interaction in EPPA is the depletion of fossil fuel resources, but the interaction effect on period 2 strategy appears to be very weak.

An interaction is also present in the climate and ocean model components. The rate of ocean uptake of carbon will gradually slow over time due to rising surface temperatures, which will cause higher period 2 carbon concentrations in the atmosphere. The slowing of ocean uptake at higher temperatures[7] becomes an interaction; higher emissions in the first period cause an increase in surface warming, which will further lower the rate of ocean uptake of CO_2 and leave higher concentrations in the atmosphere[8] (Holian, 1998). Because of the change in ocean circulation, higher period 1 emissions increase the marginal damage cost of period 2 emissions. All of the cross-period interactions in the IGSM are relatively small effects, and therefore learning does not have an appreciable influence on period 1 strategy.

Are there other possible interactions that might exist in the real world but are not represented in the IGSM? One such interaction would exist if the damage function were non-linear or had some threshold level above which marginal damages changed. The Nordhaus damage function used in this study and elsewhere, while non-linear in temperature change, is very nearly linear over the range of CO_2 concentrations resulting from a reasonable range of near-term policy choices (Pizer, 1999). Much concern about climate change is motivated by the possibility of its effect on the rate of overturning in the North Atlantic Ocean. A shift in the rate of deep-water formation, in addition to its serious climatic implications, would also alter the marginal damages of future emissions, and would constitute an inter-period interaction. Another possible effect of period 1 policy on future marginal abatement costs is *via* the rate of technological improvement. If the rate of improvement in energy efficiency and in the development of low-carbon alternative technologies can be stimulated through the presence of a price on carbon from policy, as some argue it would be (*e.g.*, Grubb, 1997), this dependence constitutes an inter-period interaction.

While each of these phenomena are argued to be important characteristics of the climate change issue, they are not represented either in the MIT IGSM or in most other climate assessment models. The main reasons for omitting them is that they are less well understood than other aspects of the system and difficult to represent in the models, and their likelihood and magnitude are highly uncertain. In the absence of these larger potential inter-period interactions,

---

[7] The solubility of CO_2 in the surface ocean layer is governed by Henry's Law, which allows the conversion between concentration and partial pressure: $[CO_2]^{sea} = a_{sol} \cdot pCO_2^{sea}$. Henry's coefficient a_{sol} has a strong dependence on temperature.

[8] The feedback mechanisms described here are distinct from an abrupt collapse or even slowdown of the thermohaline circulation. The mechanism described here is a gradual change in the rate of absorption of carbon by the surface ocean, but not in the vertical mixing between surface and deep ocean water.

studies structured as sequential decisions will never find an influence of learning on today's decision. However, it is premature to conclude from model studies that the likelihood of learning is irrelevant to the choice of near-term emission control policy. Rather, the conclusion is that the models currently used for analysis are inadequate for addressing this question.

3.5 Example of an Inter-period Interaction: Induced Technical Change

As suggested above, strong interactions can produce a learning effect. This fact can be illustrated using induced technical change as an example. EPPA, like many economic models used for climate policy studies, assumes that the rate of technological improvements is independent of the stringency of emissions policy. The appropriate model formulation for endogenous technical change and the magnitude of such an effect, if it exists, have long been debated among economists (*e.g.*, Kennedy, 1964; Samuelson, 1965; Goulder and Schneider, 1999). The model presented here is merely intended to illustrate that the influence of learning on decision changes when such a dependency is included.

We can simulate a dependency of future marginal costs on policy using the reduced-form sequential decision models, without having to modify EPPA itself. In addition to representing the uncertainty in damage costs, the reduced-form models also capture uncertainty in the abatement costs for each possible strategy. In the standard version of the sequential decision model (presented in Section 3.2), the probability distribution of the abatement costs in period 2 is exogenous and independent of the strategy in period 1. As a crude representation of endogenous technical change, we make the probability distribution of abatement costs partly a function of the period 1 strategy.

Because of the disagreement over the magnitude of an endogenous innovation effect, we parameterize the model so that the magnitude of the effect can vary from 0 (no endogenous technical change; same as original model) to 1.0 (the probability of high abatement costs is completely determined by the policy in period 1). Using this model, we want to know how the divergence between optimal period 1 strategy with and without learning will change when we include endogenous technical change. **Figure 4** shows the optimal period 1 strategies with and without learning for different probabilities of high damage cost. Here, in contrast with the results from Figure 3, there are significant regions where the strategy in period 1 depends on whether uncertainty will be resolved by period 2. The results in Figure 4 correspond to the assumption that the endogenous technical change effect is strong (0.9), and to the reference probability distributions for climate sensitivity and deep ocean uptake. Even if the endogenous technical change effect is weaker, a noticeable difference between strategies with and without learning exists.

Previous studies of the effect of learning have concluded that learning in period 2 does not affect period 1 strategy because the damage was not "irreversible enough" (Kolstad, 1996; Ulph and Ulph, 1997). This example shows that it is not that climate damages are not "irreversible enough," but rather that the IGSM does not include significant cross-period interactions. When stronger cross-period interactions are added or simulated as part of the model system, learning does have a significant effect on decision.

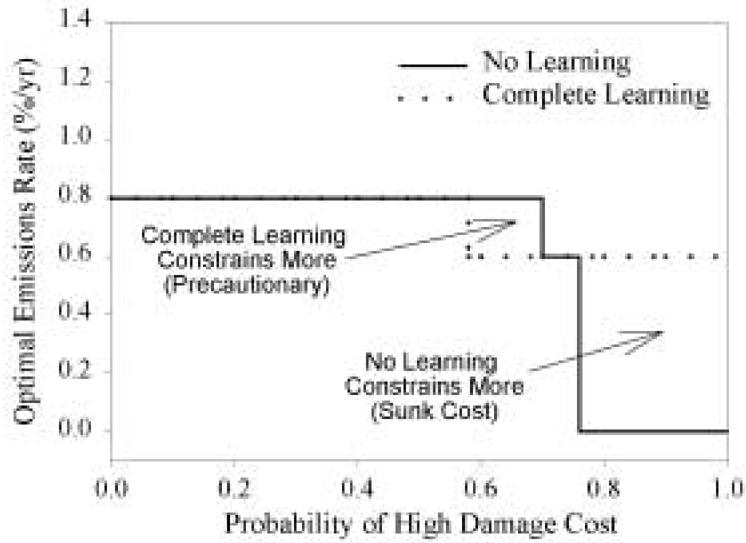

Figure 4. With endogenous technical change, learning leads to different period 1 strategy than without learning. The direction of the bias depends on the probability of high damages.

Finally, notice in Figure 4 that the direction of the bias on strategy from learning is not always in the same direction. For low expected climate damage costs (*i.e.*, low probability of the high damage cost case), the optimal period 1 strategy is the same whether or not we learn. Then for moderate expected damage costs, learning leads to more stringent emissions controls. It is worth abating more if we can learn and correct later, but not if we won't learn any more. This is a result of the skewness of the damage cost distribution as discussed in Section 2.3. The relative regret for distributions in this range is greater if we learn that high damages are the true state (the "precautionary" case). For high expected damage costs learning leads to less stringent emissions controls than if we could not learn more, because here the damage cost distribution is skewed in the direction of low damages. In this case (the "sunk cost" case), we control more if we won't learn more, but if we can learn and adjust later, we undertake less abatement. Thus the ability to learn in the future may lead to more stringent or to less stringent policy today, depending on our beliefs about the probability distributions of the damages (benefits) and of the abatement costs.

4. IMPLICATIONS FOR POLICY

This paper has explored the question of whether the ability to reduce the uncertainty about climate change in the future should lead to a delay in emissions abatement or at least to less stringent abatement. The analysis shows that whether there is an effect of learning on the first period decision depends on the existence of an interaction effect between periods. Using a climate model of intermediate complexity, it is seen that, for most parameter distributions, the optimal emissions control today is independent of whether or not learning will occur. This result can be traced to the fact that the cross-period interactions in this model are small. When an inter-period interaction is added to the model, the strategy today will depend on whether we will learn, and may lead to more or less stringent abatement depending on the relative shapes of the probability distributions of control costs and damages.

We began by considering the argument, still prevalent in climate policy discussions, that we should wait for better understanding of climate change before undertaking costly emissions abatement. Here it is shown that the ability to learn more and reduce uncertainty in the future is *not* necessarily a valid argument for delaying abatement! Learning in the future can lead to higher emissions today if we will learn, lower emissions today if we will learn, or the same emissions whether we learn or not. Which way we adjust because of learning depends on several factors including the probability distributions of the costs and benefits of emissions reductions.

What is the "act now or wait to learn" debate really about? The disagreement over the appropriate level of emissions control today is not based on differing beliefs over whether we will learn or not. The policy prescriptions are really based on differing beliefs about the expected costs of abatement and the expected climate damages. Individuals and organizations that believe that climate damages are likely to be small and that emissions abatement is likely to be costly will argue for a delay in abatement. Similarly, those who believe that climate damages are more likely to be severe and that emissions abatement may not be very costly will argue for beginning emissions abatement activities immediately. These prescriptions result directly from the perceived costs and benefits and not from considerations of the effects of reducing uncertainty in future decades.

These results have important implications for research in climate modeling. The omission of possible inter-period interactions from integrated assessment models changes the qualitative insights that emerge regarding the question what we should do today. The inclusion of possible interactions, and explicit treatment of their uncertainty should be a priority for integrated assessment modeling. Of particular importance are the 3D ocean circulation and potential thermohaline collapse, induced innovation effects of policy, and threshold effects in ecosystems damage. Of course, whether these interaction effects exist, and if so the magnitudes of the effects, are not well known. In fact, as we have shown here, it is the uncertainty in these phenomena, and not the other uncertain parameters traditionally treated, that might cause the prospect of learning to bias strategy choice under uncertainty.

Acknowledgements

I wish to thank Henry Jacoby, Gordon Kaufman, and John Reilly for invaluable advice and guidance. I gratefully acknowledge the financial support that has made this work possible from the sponsors of the Joint Program, particularly the U.S. Department of Energy (901214-HAR; DE-FG02-94ER61937; DE-FG02-93ER61713), as well as several private corporations and foundations.

REFERENCES

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

Babiker, M., J. M. Reilly, and H. D. Jacoby (1999). "The Kyoto Protocol and Developing Countries." *Energy Policy* **28**: 525-536.

Chichilnisky, G. and G. Heal (1993). "Global Environmental Risks." *Journal of Economic Perspectives* **7**(4): 65-86.

Epstein, L. G. (1980). "Decision Making and the Temporal Resolution of Uncertainty." *International Economic Review* **21**(2): 269-282.

Goulder, L.H. and S. H. Schneider (1999). "Induced Technical Change and the Attractiveness of CO_2 Abatement Policies." *Resource and Energy Economics* **21**(3-4): 211-253.

Grubb, M. (1997). "Technologies, Energy Systems and the Timing of CO_2 Emissions Abatement." *Energy Policy* **25**(2): 159-172.

Hammitt, J. K., R. A. Lempert, and M. E. Schlesinger (1992). "A Sequential-Decision Strategy for Abating Climate Change." *Nature* **357**: 315-318.

Henry, C. (1974). "Investment Decisions Under Uncertainty: The 'Irreversibility Effect.'" *American Economic Review* **64**(6): 1006-12.

Holian, G. L. (1998). *Uncertainty in Atmospheric CO_2 Concentrations from a Paramteric Uncertainty Analysis of a Global Ocean Carbon Cycle Model*. Cambridge, MA, MIT Joint Program on the Science and Policy of Global Change, Report No. 39.

Jacoby, H. D. and I. Sue Wing (1999). "Adjustment Time, Capital Malleability and Policy Cost." *The Energy Journal* **(Special Issue)** June: 73-92.

Kennedy, C. (1964). "Induced Bias in Innovation and the Theory of Distribution." *The Economic Journal* **74**(295): 541-547.

Kolstad, C. D. (1996). "Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions." *Journal of Environmental Economics and Management* **31**: 1-18.

Lempert, R. J., M. E. Schlesinger, S. C. Bankes (1996). "When We Don't Know the Costs or the Benefits: Adaptive Strategies for Abating Climate Change." *Climatic Change* **33**: 235-274.

Manne, A. S. and R. G. Richels (1995). "The Greenhouse Debate: Economic Efficiency, Burden Sharing and Hedging Strategies." *The Energy Journal* **16**(4): 1-37.

Manne, A. S. and R. G. Richels (1992). *Buying greenhouse insurance: the economic costs of carbon dioxide emission limits*. Cambridge, Mass., MIT Press.

Miller, J. R. and F. Lad (1984). "Flexibility, Learning, and Irreversibility in Environmental Decisions: A Bayesian Approach." *Journal of Environmental Economics and Management* **11**(2): 161-172.

Morgan, M. G. and M. Henrion (1990). *Uncertainty: A Guide To Dealing With Uncertainty In Quantitative Risk And Policy Analysis*. Cambridge; New York, Cambridge University Press.

Morgan, M. G. and D. W. Keith (1995). "Subjective Judgments by Climate Experts." *Environmental Science and Technology* **29**(10): 468A-476A.

Nordhaus, W. D. (1994). *Managing The Global Commons: The Economics Of Climate Change*. Cambridge, Mass., MIT Press.

Nordhaus, W. D. (1994b). "Expert Opinion on Climatic Change." *American Scientist* **82**(January): 45-51.

Nordhaus, W. D. and D. Popp (1997). "What is the Value of Scientific Knowledge? An Application to Global Warming Using the PRICE Model." *The Energy Journal* **18**(1): 1-45.

Peck, S. C. and T. J. Teisberg (1992). "CETA: A Model for Carbon Emissions Trajectory Assessment." *The Energy Journal* **13**(1): 55-77.

Pindyck, R. S. (1999). *Irreversibilities and the Timing of Environmental Policy*. Working Paper 99-005, Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology, Cambridge, MA.

Pindyck, R. S. (1991). "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature* **29**(3): 1110-48.

Pizer, W. A. (1999). "The optimal choice of climate change policy in the presence of uncertainty." *Resource and Energy Economics* **21**(3,4): 255-287.

Prinn, R., H. Jacoby, *et al.* (1999). "Integrated Global System Model for Climate Policy Assessment: Feedbacks and Sensitivity Studies." *Climatic Change* **41**(3/4): 469-546.

Risbey, J. S., M. D. Handel, and P. H. Stone (1991a). "Should We Delay Responses to the Greenhouse Issue?" *EOS, Transactions* **72**(53): 593.

Risbey, J. S., M. D. Handel, and P. H. Stone (1991b). "Do We Know What Difference a Delay Makes?" *EOS, Transactions* **72**(53): 596-597.

Roughgarden, T. and S. H. Schneider (1999). "Climate Change Policy: Quantifying Uncertainties for Damages and Optimal Carbon Taxes." *Energy Policy* **27**: 415-429.

Samuelson, P. A. (1965). "A Theory of Induced Innovation Along Kennedy-Weisacker Lines." *The Review of Economics and Statistics* **47**(4): 343-356.

Schlesinger, M. E. and X. Jiang (1991a). "A Phased-in Approach to Greenhouse-Gas-Induced Climatic Change." *EOS, Transactions* **72**(53): 593,596.

Schlesinger, M. E. and X. Jiang (1991b). "Climatic Responses to Increasing Greenhouse Gases." *EOS, Transactions* **72**(53): 597.

Sokolov, A. P. and P. H. Stone (1998). "A Flexible Climate Model for use in Integrated Assessments." *Climate Dynamics* **14**: 291-303.

Stevens, W. K. (1997). On Global Warming, Some in Industry Are Now Yielding. New York Times. New York.

Tatang, M. A., W. Pan, *et al.* (1997). "An efficient method for parametric uncertainty analysis of numerical geophysical models." *Journal of Geophysical Research* **102**(D18): 21.

Ulph, A. and D. Ulph (1997). "Global Warming, Irreversibility and Learning." *Economic Journal* **107**(442): 636-650.

United Nations (1992). "Framework Convention on Climate Change." *International Legal Materials* **31**: 849-873.

United Nations (1997). Kyoto Protocol to the United Nations Framework Convention on Climate Change. FCCC/CP/1997/L.7/Add.1. Bonn.

Viscusi, W. K. and R. Zeckhauser (1976). "Environmental Policy Choice under Uncertainty." *Journal of Environmental and Economic Management* **3**: 97-112.

Webster, M. D. (2000). *Uncertainty and Learning in Sequential Decision-Making: The Case of Climate Policy*. Ph.D. Thesis, Massachusetts Institute of Technology, Technology, Management and Policy.

Webster, M. D. and A. P. Sokolov (2000). "A Methodology for Quantifying Uncertainty in Climate Projections." *Climatic Change* **46**(4): 417-446