

Incorporating Path Dependency into Decision-Analytic Methods: An Application to Global Climate-Change Policy

Mort Webster

Earth Atmosphere and Planetary Sciences Department, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139, mort@mit.edu

Climate policy decisions are necessarily sequential decisions over time under uncertainty, given the magnitude of uncertainty in both economic and scientific processes, the decades-to-centuries time scale of the phenomenon, and the ability to reduce uncertainty and revise decisions along the way. Thus, an appropriate choice of analytical method is decision analysis. However, applying decision analysis in the context of idealized government decision makers over a century raises the question of how to deal with the fact that political systems tend to exhibit path dependency, a force that makes large policy shifts difficult and rare, and limits most decisions to small incremental changes. This paper explores the effect of considering path dependency in an application of decision analysis to climate-change policy decisions, presenting two alternative methods for modeling path dependency. I demonstrate that consideration of path dependence in the context of climate policy justifies greater near-term emissions reductions. The more general result of path-dependency is to shift the near-term strategy towards a more moderate hedging strategy, because drastic shifts later will be difficult.

Key words: climate policy; decision analysis; path dependency

History: Received on December 21, 2007. Accepted on May 5, 2008, after 2 revisions.

1. Introduction

Applications of decision analysis to public policy decisions, especially policies that span long time horizons, face an additional challenge as compared with individual decision. The challenge is the tendency of policies and government programs to be “sticky;” i.e., it is difficult to make more than small incremental changes at any one decision point. This tendency of policies to exhibit path dependency has long been noted by political scientists in a variety of contexts. In this paper, I present a methodology to account for path dependency in political decisions in an example application to climate-change policy and greenhouse gas reductions.

Formulating a policy response to the threat of global climate change is one of the most complex public policy challenges of our time. One troubling characteristic is the enormous uncertainty involved, both in the magnitude of future climate change, and therefore the value of avoiding it, and in the costs of reducing emissions. The long time scales of the

climate system, decades to centuries, add another dimension to the policy dilemma. Given the stock nature of greenhouse gases, which build slowly over time and have a long lifetime, should we delay mitigation activities until some of the uncertainties are reduced?

Issues of uncertainty and whether to delay mitigation are central to the climate policy debate, particularly in the United States. The current U.S. policy is one of strictly voluntary emissions targets; firms who wish to reduce their greenhouse gas emissions may do so. The literature on public goods problems predicts that in this type of situation, the level of the good provided (emissions reductions) will be far below the optimal level, perhaps zero (Weimer and Vining 2005). The Bush administration argues that legally mandated emissions reductions should be delayed for at least the next decade (Bush 2002). One of the reasons cited is the uncertainty in future climate change; mandatory restrictions are too costly given the incomplete state of scientific knowledge of

the causes of, and solutions to, global climate change (Bush 2001). Alternative regulatory proposals of various levels of effort continue to be debated, including the Liebermann-Warner (Liebermann 2007) and the Bingaman-Specter (Bingaman 2007) proposals in the U.S. Senate, and the National Commission on Energy Policy recommendation (NCEP 2004). However, to date, no consensus on any alternative has emerged.

The state of international climate policy is more similar to that in the United States than it may first appear. Although the Kyoto Protocol to the United Nations Framework Convention on Climate Change has entered into force with binding targets for many developed nations, the United States and Australia, as well as developing countries, are not bound by targets. Further, it is not yet clear whether significant reductions will result in the countries with Kyoto commitments (Reilly and Paltsev 2005, *The Economist* 2006). Nor is it clear what further reductions, if any, the international community is willing to undertake post-2012 (Revkin 2005). Thus, the central question for near-term climate policy, both in the United States and abroad, is whether or not regulations of greenhouse gas emissions can be delayed for another decade or whether some level of mitigation effort is required now.

When irreversibilities exist in the presence of uncertainty, delay is not necessarily optimal. In the classic works of Arrow and Fisher (1974) and Henry (1974), they demonstrated that there is an additional value, called the quasi-option value, to preserving some of a resource under uncertainty and irreversibility. However, in the climate-change problem, there are irreversibilities in both directions: Changes to the physical climate system are irreversible, but so are changes in the capital stock to lower emissions, as identified by previous studies (Kolstad 1996, Ulph and Ulph 1997, Ha-Duong 1998, Webster 2002). Numerous analyses have used dynamic optimization models to examine the optimal levels of greenhouse gas mitigation, both under conditions of certainty and uncertainty (Hammit et al. 1992, Manne and Richels 1995, Nordhaus 1994a, Nordhaus and Popp 1996, Kolstad 1996, Ulph and Ulph 1997, Ha-Duong 1998, Valverde et al. 1999, Webster 2002). These studies have generally demonstrated an optimal near-term

climate policy of very little or no emissions reductions. This is because after resolution of uncertainty, if it is revealed that climate change is very serious, then more stringent policies can be pursued in later decision periods. Thus, quantitative climate policy studies have appeared to lend support for delayed mitigation.

However, there is a critical element that is missing from both the policy debate and from the formal models of climate policy: path dependency. Political scientists have long noted the tendency of political systems to exhibit path dependency, and have used this feature to explain a number of political outcomes, such as European party systems (Lipset and Rokkan 1967) and the comparative development of healthcare systems (Hacker 1998). The idea of path dependency is that once a particular course of action has been chosen, it becomes increasingly difficult over time to reverse that course (Pierson 2000, Sewell 1996, Levi 1997). Policies tend to exhibit lock-in, and although a legislature might from time to time create a new bureaucratic agency, it is exceedingly difficult to eliminate one. Path dependency not only affects the creation/elimination of programs, but may also account for the difficulty in adjusting stringency of policies such as tax rates (Kaplow and Shavell 2002) or sulfur emissions cap levels (Ellerman et al. 2000). Path dependency has been recognized and modeled, either as increasing returns to production or as transaction costs, in a variety of economic policy applications as well, including money supply adjustments (Dixit 1991), long-run growth equilibria of firms (Altman 2000), and optimal resource extraction (Gerlagh and Keyzer 2004). Zhao and Kling (2003) modeled path dependency as transaction costs to show that policy persistence may be a rational forward-looking response.

A large-scale international policy issue such as climate change is especially vulnerable to path dependencies. If significant global emissions reductions are required in the long-run, this will be an extremely difficult problem to coordinate across nations. One analog that one might consider is the development of the World Trade Organization (WTO). In 1947, the General Agreement on Tariffs and Trade (GATT) was signed with the goal of eliminating all barriers to international trade. After 50 years, there has been much progress in the development of the institutional

capability, but complete free trade is still far from realized. A global effort to reduce greenhouse gases on the scale required to stabilize concentrations would be at least as ambitious as the WTO. Further delay in the development of the institutional capacity will limit the ability to respond if climate change turns out to be serious.

The problem of path dependency has been modeled in some applications of dynamic programming (e.g., Bertsimas et al. 2001 using a jump-diffusion process). However, in climate policy studies using these techniques, and in most decision analysis applications to policy, path dependency has been notably missing from the models used. Climate policy optimization models typically assume that some fraction of baseline emissions can be reduced in each period, ranging from none to nearly 100%. However, the range of reductions considered in any period is independent of any choices made in previous periods.

The question posed in this study is: Does accounting for the path dependency in political systems change the first-period (today) optimal choice from a sequential decision model of climate policy? If it does, then this would argue for a more aggressive hedging strategy with greater emissions reductions for near-term climate policy. This action would allow for greater flexibility if significant reductions are required later in the century. The primary contributions of this study are conceptual and policy prescriptive. I develop here a simple illustrative model of path dependency to demonstrate its importance in near-term policy considerations.

Section 2 describes the model that is used to project climate outcomes and costs of policies, and to find least-cost emissions paths over time under uncertainty. I use this model as a context in which to explore the implications of path dependency for climate policy. I will compare the results of decision models with differing degrees of path dependency in §3. The final section discusses the implications both for climate policy and for research.

2. Experimental Design

There is a wide spectrum of models that can be used to project the impacts of greenhouse gas emissions and resulting climate change as well as the economic costs of constraining those emissions. These

range from very simple approximations to very large sophisticated models that require weeks on a supercomputer to simulate. The advantage of the more complex models is that they represent many of the nonlinearities and complexities that make climate change a cause for concern. On the other hand, the requirements of solving a dynamic optimization under uncertainty require some simplification to make the analysis feasible. The approach used here is to obtain emissions and cost impacts from a relatively detailed computable general equilibrium model of the global economy, combined with the climate impacts of the emissions obtained from a reduced-form model calibrated to a climate model of intermediate complexity. The resulting costs and climate impacts are then embedded within a decision tree framework, which is used to solve the intertemporal decision problem under uncertainty.

2.1. The Decision Model

The decision model is a two-period model, with uncertainty in several key parameters, and learning between the first and second decisions. The objective of the decision maker is to minimize the present value of the stream of total costs (TC), which are the sum of abatement costs (AC) from reducing emissions below their baseline levels and damage costs (DC) from climate change,

$$TC = \sum_{t=1}^2 (AC(t) + DC(t)). \quad (1)$$

The climate damages are determined by δ , which is a function of temperature change $\Delta T(t)$ and a damage valuation (DV) elasticity parameter α , following Nordhaus (1994a):

$$DC(t) = \delta(\Delta T(t), DV). \quad (2)$$

The function α gives the lost welfare in each period as a function of the divergence between the carbon emissions $G(t)$ and the emissions with no policy $G^*(t)$.

$$AC(t) = \alpha(G(t), G^*(t)). \quad (3)$$

The temperature change in Equation (2) is a function of greenhouse gas emissions G and of three uncertain climate parameters, CS , Kv , and Fa :

$$\Delta T(t) = \tau(G(t), CS, Kv, Fa). \quad (4)$$

The climate sensitivity CS is the equilibrium temperature change resulting from a doubling of CO_2 concentrations, and represents aggregate feedbacks in the atmosphere; the deep ocean heat uptake Kv represents the mixing of heat between the surface layer and deeper layers of the ocean; and the aerosol forcing strength Fa represents the magnitude of the negative radiative forcing of a unit loading of sulfate aerosol in the atmosphere.

Carbon emissions G in Equations (3) and (4) are determined by the function ε , which is a function of the emissions reduction rate that is chosen (the control variables) in each period ERR_i and by the uncertain labor productivity growth (LPG) rate. Emissions in period 1 depend on the policy choice in period 1, but emissions in the second period depend on the policy in both periods:

$$\begin{aligned} G_1 &= \varepsilon(ERR_1, LPG) \\ G_2 &= \varepsilon(ERR_1, ERR_2, LPG). \end{aligned} \quad (5)$$

Note that the no-policy emissions G^* is simply:

$$G^* = \varepsilon(0, 0, LPG). \quad (6)$$


We assume that the second-period decision is made conditional upon information z , which the decision maker receives prior to the second decision. Information z indicates the true value of the uncertain parameters. Thus, the decision-maker's problem is to choose emission reduction rates in each period, such that the expected present value of total costs is minimized:

$$\text{MIN}_{ERR_1, ERR_2} = E\{TC | z\} \quad \text{s.t. Equations (1)–(6),} \quad (7)$$

where $E\{\cdot\}$ is the expectation with respect to uncertain parameters LPG, CS, Kv, Fa , and DV .

We illustrate the model components in Figure 1. The abatement cost function in Equation (3) and the emissions function in Equations (5) and (6) are implemented in the Emissions Projections and Policy Analysis (EPPA) model (Paltsev et al. 2005), a recursive-dynamic computable general equilibrium model, consisting (in the calculation applied here) of 12 geopolitical regions linked by international trade, 10 production sectors in each region, and 4 consumption sectors. It has been used for numerous analyses

Figure 1 Schematic of Model Components



of climate policy (e.g., Babiker et al. 2002, Jacoby et al. 1997, McFarland et al. 2004, Reilly et al. 2002). The temperature-change function in Equation (4) is based on the MIT climate model (Prinn et al. 1999, Sokolov and Stone 1998), a two-dimensional (zonal averaged) representation of the atmosphere, ocean, and terrestrial biosphere. The climate model includes parameterizations of all the main physical atmospheric processes, and is capable of reproducing many of the nonlinear interactions simulated by atmospheric general circulation models (GCMs). Because of the computational cost of simulating the full climate model for all possible policies and uncertainties in this study, I use a reduced-form model fit to simulations of the climate model using the method described in Tatang et al. (1997) and described in detail in Webster et al. (2003), and Webster and Sokolov (2000). The damage function in Equation (2) is based on Nordhaus (1994a) and is of the specific form:

$$\delta(t) = DV[\Delta T(t)]^2 \quad (8)$$

where $\delta(t)$ is the fraction of world product lost due to climate damages in year t , DV is the percentage loss from a doubling of CO_2 concentrations, and $\Delta T(t)$ is the increase in global mean temperature from preindustrial levels. Abatement and damage costs are discounted sums over time, using a discount rate of 3%.

2.2. Decision Periods, Strategy Space, and Distributions of Uncertain Parameters

To focus on the dynamics of optimal decision in the presence of path dependency, I simplify the model to two time periods: one from 2010–2030 and the second from 2030–2100. The assumption of a two-period model follows a long tradition of using simple two-period models to develop intuition, including Arrow and Fisher (1974), Manne and Richels (1995), Hammitt

Table 1 Strategy Choices in Two-Period Model: Reduction Below Unconstrained Emission Growth Rate

Decision period	Strategy variable	Years	Strategy choices available (reduction below unconstrained rate) (%)
1	Policy 2010	2010–2029	0, 2, 4, 6, 8, 10
2	Policy 2030	2030–2100	0, 1, 2, 3, 4, 5

et al. (1992), Yohe et al. (2004), and others. A model with more than two time periods would not change the fundamental intuition for the first-period optimal strategy. To demonstrate this last point, in §3.4 I test the sensitivity to the number of periods, using a three-period version of the model presented here.

Another simplification is the focus on global outcomes. In this paper, the goal is to explore the implications of path dependency for overall level of abatement, and avoid questions of relative burden sharing among nations. I assume global trading of emissions permits between all countries, and only examine the total global losses. To add a sense of reality, only half of the emissions reductions described below are applied to developing (nonAnnex I) nations between 2010 and 2040. After 2040, all policies apply equally. Other assumptions about the relative participation of developing countries would not change the qualitative results of this analysis.

Strategies in each period are defined as the reduction required in the rate of growth of carbon emissions, relative to the unconstrained case. Thus, these policies will vary with the (uncertain) rate of economic growth. Thus, 0% means no emissions constraints at all, and 5% means a 5 percentage point reduction in the CO₂ growth rate over that five-year period, relative to the reference rate of emissions growth. For example, if a region's emissions grow at 5% under no policy, then the 5% policy would result in emissions stabilization (0% growth allowed).

Smaller rates of reduction would result in slowed growth of CO₂ emissions, whereas larger rates would actually reduce global emissions over time.

The set of available strategies for each decision period are given below in Table 1. The emissions from this strategy set define an envelope between a no-policy case and stringent reductions over the century that nearly stabilize carbon concentrations at 550 ppm. To put these policy choices in more familiar terms, Table 2 lists the impacts of each possible first-period strategy by 2030 for the median productivity growth case, and the initial carbon price in 2010. The numbers in Table 2 are the EPPA model results under each of these policy constraints, and indicate the values for the year 2030 as an example. Values under these policies in other years during the first period (2010–2030) would vary because the economy is growing, but the proportional effects will be similar to those in 2030. The reason that the magnitudes of period 1 reduction rates are larger than those in period 2 is a function of the relative difference in the time horizon in each period (period 1 is 20 years, period 2 is 70 years) and of the range of reductions needed to span emissions paths from no policy to 550 ppm CO₂ stabilization. In the presentation of results (§3), we describe the period 1 strategies in terms of the 2010 carbon tax level, because this is a more intuitive description for stakeholders in the climate policy debate.

Based on previous work (Webster et al. 2003, 2002), the model includes five uncertain parameters that have the greatest impact on damage costs: *LPG*, *CS*, *Kv*, *Fa*, and *DV*. The three uncertain climate parameters, *CS*, *Kv*, and *Fa*, are combined for each possible emissions path by performing a Monte Carlo simulation of 10,000 trials on the reduced-form climate model. The total resulting uncertainty in temperature

Table 2 Impacts of Period 1 Strategy Choice in 2030 (Median Growth Case)

Reduction rate/5yrs (%)	CO ₂ (GtC)	Chg CO ₂ (from BAU) (%)	Carbon price (2010) \$/ton C	Carbon price (2030) \$/ton C	Consumption (billion \$)	Chg Cons (%)
0	12.5		0	0.0	5,094.2	
2	11.8	–5	5.0	19.3	5,091.8	–0.05
4	11.1	–11	12.0	45.1	5,087.9	–0.12
6	10.5	–16	20.0	79.7	5,082.1	–0.24
8	9.9	–20	27.0	125.1	5,073.3	–0.41
10	9.4	–25	36.0	181.6	5,060.8	–0.66

Table 3 Distributions for Uncertain Quantities

	Branch 1 ($P = 0.185$)	Branch 2 ($P = 0.63$)	Branch 3 ($P = 0.185$)
Labor productivity growth rate (relative to reference rates)	0.8	1.0	1.2
Temperature change (degrees C)	5th percentile	Median	95th percentile
Damage cost coefficient (%)	0.02	0.04	0.16

change is then summarized by a three-point Tukey-Pearson approximation (Keefer and Bodily 1983) using the 5th, 50th, and 95th percentiles. The joint probability distributions of the climate parameters were obtained from a climate detection study of observations of 20th-century climate (Forest et al. 2002, Webster et al. 2003). The uncertainties in labor productivity and damage valuation are also represented in the decision tree with three-point discrete approximations, as described in Table 3. The distribution of labor productivity growth is based on Webster et al. (2002). The distribution for the damage valuation is taken from Roughgarden and Schneider (1999), based on the assessment by Nordhaus (1994b).

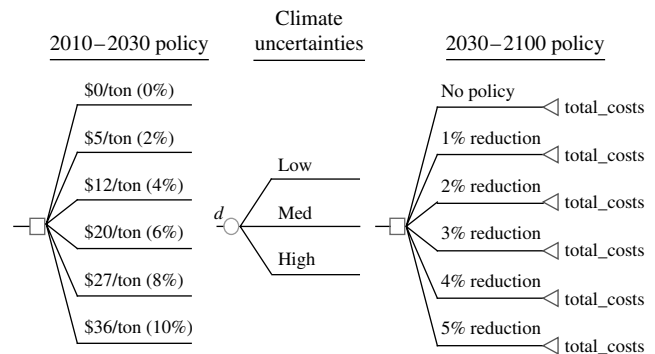
The concern over the presence of path dependency is most relevant in the case where uncertainty is resolved. If uncertainty is known to be irresolvable, the optimal path of emissions over the next century is relatively smooth, and no large shifts in policy are expected within any decade. However, if uncertainty is expected to be reduced, one of the justifications for delaying mitigation policy, a sudden increase or decrease in the stringency of policy may be required. Thus, I focus on a polar case where uncertainty is resolved completely in 2030, consistent with the tradition of sequential climate policy decision models (Hammitt et al. 1992, Manne and Richels 1995, Nordhaus and Popp 1997, Ulph and Ulph 1997, Ha-Duong 1998, Valverde et al. 1999, Webster 2002).

This two-period decision model is shown in Figure 2 as a fragment of a decision tree. In the first decision period, the policy for 2010–2030 is chosen, then climate uncertainties are resolved, and for every possible resolution, the policy for 2030–2100 is then chosen.

2.3. Modeling Path Dependency

A challenge in exploring this issue is *how* to represent path dependency in dynamic optimization mod-

Figure 2 Decision Tree Fragment for Standard Decision Model with No Path Dependence



els of the type used here. Ideally, this feature would be represented somehow in the underlying representations of the costs and benefits of each decision path, as calculated by the economic and climate models. Because path dependency is not embedded in the models, however, the goal here is to add a relatively simple adjustment to the decision model that has the desired effect and that makes sensitivity analysis straightforward.

There are several possible representations of path dependency. One straightforward method is to model an additional component of total cost that is an increasing function of the distance between the chosen policy and its previous level. A second alternative is to add constraints that prevent a policy choice that is “too far” from the previous policy, essentially limiting each decision point to a maximum incremental shift of stringency in either direction. A third approach is to model policy decisions in the future not as a decision, but rather as an uncertainty that the current decision maker cannot control. Finally, one could model path dependency as a time lag between a decision and the implementation. I discuss the relative advantages of each approach briefly, and then describe the methodology for two of the above approaches that are modeled below.

The time lag approach is the least flexible in terms of capturing the behavior that leads to small rather than large changes at a point in time. Adding a time delay does not by itself constrain the size of the policy shift. A constraint on feasible options comes closer. The disadvantages of the constraint approach are (1) the arbitrariness in choosing which policies are

not feasible given the previous period; and (2) the discontinuous nature of allowing some shift with no penalty, but making the penalty infinite beyond a given magnitude. In this paper, we implement the other two approaches: the cost function approach and the probabilistic approach. Of the two, the cost approach is simpler. The primary disadvantage is the difficulty in assessing an appropriate magnitude for the cost penalty. The probabilistic approach we present is more complex, but it captures the fact that the current decision makers are not the ones who will make the future decisions.

To model path dependency as a cost function that increases in the distance between the current and previous policies, we modify the calculation of total costs (Equation (1)) to include the additional cost term:

$$TC = \sum_{t=1}^2 (AC(t) + DC(t)) + \lambda(ERR_2 - ERR_1)^2. \quad (9)$$

The parameter λ represents the magnitude of the cost penalty, where higher values will induce greater path dependency.

In the second approach I assume that because the relevant decision makers are choosing for the present only and have no control over future decisions, those future choices can be modeled as uncertainties rather than decisions. Within this framework, path dependency is modeled by assigning a probability to a future decision that is proportional to its difference from the earlier decision.

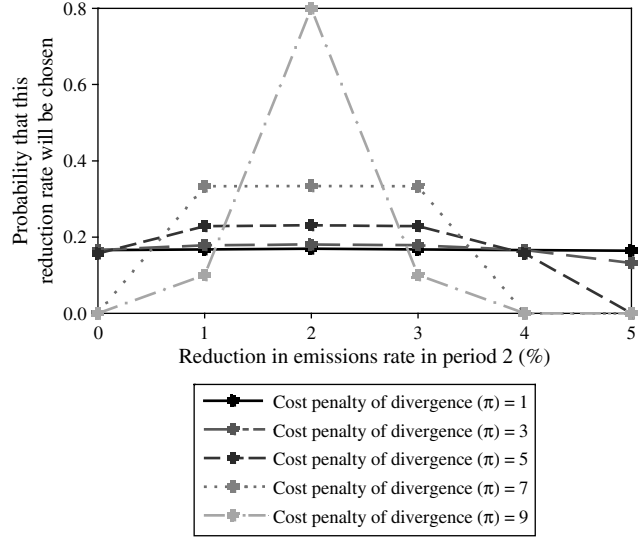
To begin, we can represent period 2 decisions with a probability distribution by defining the probability of choosing the optimal strategy as:

$$f_o(ERR_2 | z) = \begin{cases} 1 & \text{if } ERR_2 = \arg \min(TC) \\ 0 & \text{otherwise} \end{cases}. \quad (10)$$

In words, this distribution defines the probability of a period 2 strategy to be one if it is the optimal strategy in the standard two-period model, and otherwise it is zero. Using this distribution over period 2 strategies produces an equivalent decision model to the original two-period model described above.

Next, we introduce a probabilistic representation of path dependency by letting the probability be decreasing in the distance between the stringency of the current policy and the stringency of the previous

Figure 3 Sample Path-Dependent Probability Distributions for Period 2 Policies Given Period 1 Choice of $ERR_1 = 2\%$ (\$12/ton C in 2010)



policy. Specifically, we define a probability distribution over the emissions reduction rate in period 2, ERR_2 . We define the path dependent probability density function of the second-period policy as:

$$f_p(ERR_2) \propto C - (ERR_2 - ERR_1)^\pi, \quad (11)$$

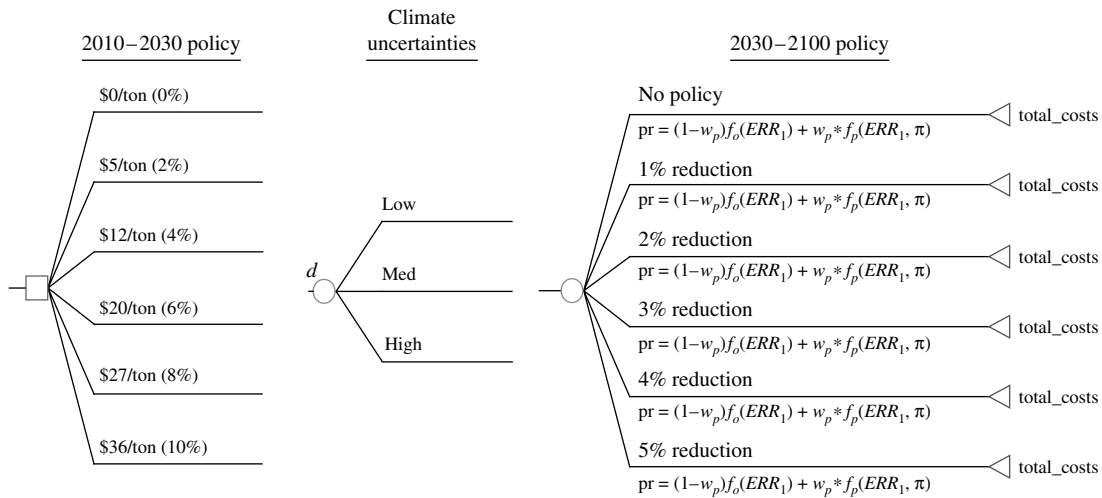
where C is an arbitrary constant and the distribution is normalized to integrate to one. The parameter π is the penalty for divergence from the previous period, thus representing the severity of the path dependent effect. I consider values of π from 1 to 10. Figure 3 graphically shows an example of the probability distributions over period 2 policies, assuming a period 1 choice of 4% or \$12/ton carbon tax, for several possible values of π .

The overall probability distribution over future policies is a weighted sum of these two influences: the pull of the optimal¹ policy and the inertia of politics and institutions. This is implemented as a linear mixture of the two probability distributions defined above:

$$\Pr(ERR_2) = (1 - w_p)f_o(ERR_2) + w_p f_p(ERR_2), \quad (12)$$

¹ In terms of minimizing the sum of abatement costs and climate damage costs, but ignoring all other factors.

Figure 4 Decision Tree Fragment for Standard Decision Model with Probabilistic Path Dependence



Note. The second-period decision now becomes a chance node, and the probability is the weighted average of the being the purely optimal decision and distance from the previous policy.

where w_p is the relative weight on path dependency in determining the period 2 strategy, and ranges from 0 to 1.

Path dependency is therefore represented by changing the second-period decision node to a chance node, and the probability of every branch from that node is defined as in Equation (12). This formulation gives two parameters that can be varied to explore the full range of path dependency in the system: w_p , the relative influence of path dependency on the decision, and π , the penalty for deviating from the previous decision. Note that when w_p is zero, the decision model will produce identical results to the standard two-period decision model without path dependency. The intuition is that π represents how much more difficult it is to diverge from the current policy trajectory the further the divergence, whereas w_p represents the relative balancing in future decision making between purely climate-oriented cost effectiveness and the drag of political inertia.

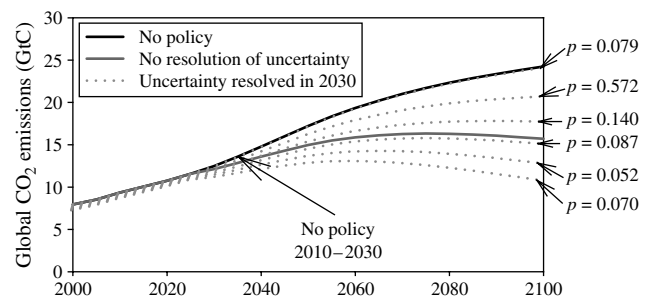
Figure 4 shows the path dependent version as a fragment of the decision tree. Very simply, we convert the period 2 decision node from Figure 2 to a chance node. The probability of any branch emanating from the period 2 policy chance node is determined by Equations (10)–(12) above.

3. Results

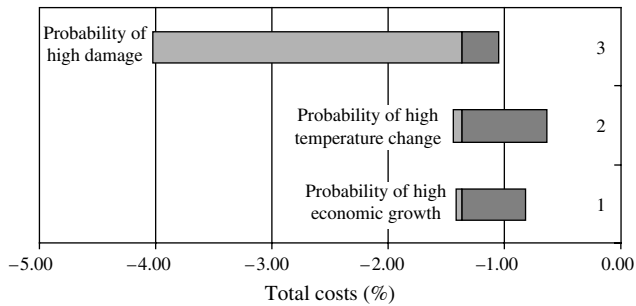
3.1. Optimal Decision with No Path Dependency

Before showing the effect of incorporating path dependency on the optimal sequential model, we begin with the solution to the standard two-period model, which we will refer to here as the no-path-dependence version (see Figure 2). Figure 5 shows the optimal global path of CO₂ emissions for two variations on the model without path dependency: one in which uncertainty is completely resolved between the two decision periods and one in which the uncer-

Figure 5 Solution for Standard Model (No Path Dependence) and Reference Probability Distributions



Notes. Solid line indicates emissions with no policy; dashed line indicates the optimal emissions path if uncertainties are not resolved before 2100; dotted lines indicate the optimal emissions paths if uncertainty is resolved in 2030. The associated probabilities p give the probability that the resolution of uncertainty in 2030 will result in that path being optimal.

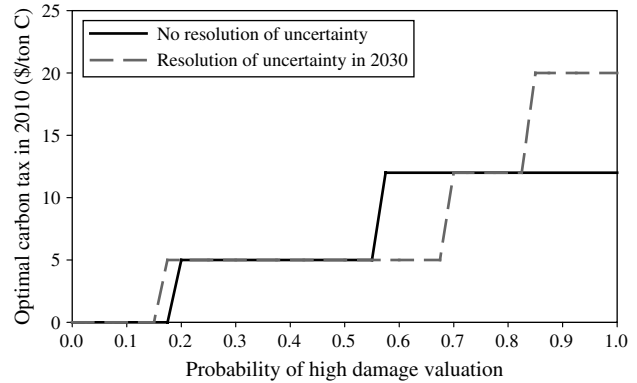
Figure 6 Relative Influence of Uncertainties on Total Costs

Notes. Each probability of the “high” case was varied from 0.0 to 1.0, and remaining probability is split 80%/20% between the median and low cases, respectively. Dark shading in bars indicates that the optimal decision is No Policy (\$0/ton CO₂) and light shading indicates that optimal decision is \$5/ton CO₂.

tainty is not resolved at all until after both decisions have been made. The important feature to notice is that the emissions between 2010 and 2030 (the first decision period) do not diverge from the no-policy emissions path. In other words, the optimal first-period decision is to undertake no abatement. The emissions paths from 2030 to 2100 demonstrate that in the second decision period, it will almost always be optimal to reduce emissions from the no-policy case; how much of a reduction is optimal depends on what one learns.

This result by itself is not convincing because of disagreement over what the probability distributions of the uncertain parameters are. The distribution for climate-damage valuation is particularly contentious (Nordhaus 1994b, Roughgarden and Schneider 1999). The damage valuation uncertainty also has the strongest impact on the optimal policy choice. Figure 6 shows a tornado diagram for this decision model, which ranks the relative impact of the three uncertainties in this model: economic productivity growth, physical climate uncertainties, and damage valuation.

To further explore the conditions under which first-period emissions reductions would be optimal, we subject the three-point discrete approximations of the probability distributions to sensitivity testing, by varying the probability of the high value, while fixing the probability ratio of median to low damage value. Figure 7 shows the optimal period 1 strategy, in terms

Figure 7 Sensitivity Analysis for the Optimal First-Period Decision for the Standard Model (No Path Dependence) as the Probability of High Damage Valuation Varies from 0.0 to 1.0

Note. Solid line shows optimal strategy if uncertainty is never resolved, and dashed lines show optimal strategy if uncertainty is resolved in 2030.

of the carbon tax in 2010,² as a function of the probability of the high climate-damage value. If the probability that climate damage has significant economic impacts is below 0.2, it is optimal to have no emissions reductions. For probabilities of high damage between 0.2 and 0.7, the optimal carbon tax is \$5/ton, and the optimal tax is increasing in the probability of high climate damage. However, these higher probabilities (>0.4) that justify higher carbon taxes are significantly higher than the assessments of most experts (Nordhaus 1994b). Note that for different assumptions about the value of climate damage, the ability to learn in 2030 may result in a more or a less stringent policy than one with no learning. As demonstrated elsewhere (Webster 2002), this is the result of opposing irreversibility effects: the capital stock in the economy and change to the climate system. Depending on the shape of the distribution over uncertain net costs change, the regret over one irreversible effect will dominate regret over the other, leading to hedging against the worse outcome.

One useful quantity for a decision problem under uncertainty is the expected value of perfect information (EVPI). For the standard model with resolution

² The results for optimal period 1 strategies will be given in terms of the equivalent carbon tax in 2010, as listed in Table 2, in order to frame the results in terms of quantities that may be more familiar to the reader. Carbon taxes listed here are in dollars per ton of carbon; an alternative commonly used is per ton of CO₂, which can be converted to \$/ton carbon by multiplying by 12/44.

of uncertainty in 2030 and reference probability distributions, the EVPI is 0.19%, or \$79 billion discounted to 2000 at 5%. However, the EVPI is a comparison between the decision in which we choose the near-term policy under uncertainty and one in which we choose near-term policy under perfect knowledge. Thus, although EVPI indicates that the value of reducing uncertainty is substantial, it does not give us any guidance into whether further delay of mitigation is good policy or not.


As discussed in the introduction, this result is consistent with other studies in the literature. Several characteristics of this problem cause the economically efficient solution to delay the bulk of abatement activities for several decades, including the long lifetime of CO₂ and the inertia in the climate system, the technological improvements that occur over time, and discounting over time. What we wish to explore here is whether there is an additional effect from the very structure of the decision model itself.

3.2. Optimal Decision with Path Dependency Influenced by Period 1 Decision (Cost Approach)

Next, I present results from the modified version with path dependency, as described in §2.3. First I will compare the behavior of the cost function approach to path dependency with the results above. In the following section, I will present the results from the probabilistic approach to path dependency. In the cost function approach, the strength of the path dependent effect will increase with the cost penalty parameter λ . We are primarily interested in whether including the effect of path dependency in the model alters the optimal first-period strategy, and if so, how strong the effect is.

Figure 8 reproduces the sensitivity of optimal first-period strategy from Figure 7 (for the case with resolution of uncertainty in 2030), and includes the results for three different levels of cost penalty. If the optimal decision in period 2 is further from the previous decision, the costs are increased, thereby constraining large shifts in stringency. In general, the presence of path dependency results in the optimal period 1 policy becoming more stringent, justifying a higher carbon price in the near term. However, a greater cost penalty does not necessarily increase the shift towards

Figure 8 Optimal First-Period Strategy as a Function of the Probability of High Climate Damages

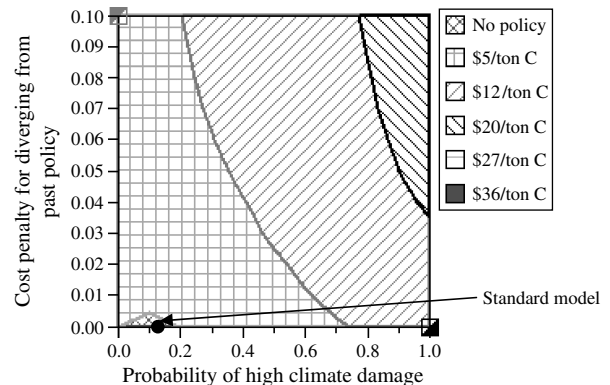


Note. Solid line shows solution with no path dependency, and the dashed lines show three different levels for the cost penalty of diverging from previous policy.

higher carbon taxes; it depends on the probability of high climate damages.

A more complete picture of the effect of path dependency is shown in Figure 9. This graph displays a two-way sensitivity of varying both the probability of high damage and the cost penalty on policy shifts; the shading pattern indicates the optimal period 1 policy choice. The block dot indicates the standard version of the model, with no path dependency and a probability of high damage of 0.185. Under the assumptions of the standard model without path dependency, a zero carbon price was justifi-

Figure 9 Sensitivity of Optimal Period 1 Decision to the Probability of High Damage and to the Cost Penalty of Divergent Policy



Note. Black dot indicates the reference version of the model with no path dependency.

fied in the near term. However, if even a very slight influence of path dependence exists, the optimal policy shifts to a carbon price of \$5/ton in 2010.

3.3. Optimal Decision with Path Dependency Influenced by Period 1 Decision (Probabilistic Approach)

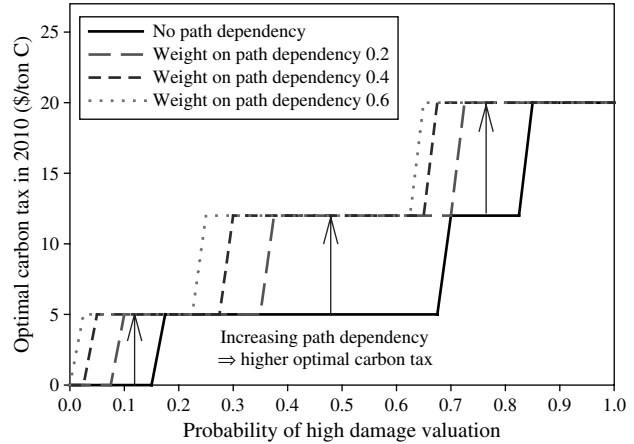
As described in §2.3, an alternative approach to modeling path dependency is to represent the period 2 decision as an uncertainty. This alternative approach might be more appropriate for very long time horizons, such as in this application to climate-change policy. In this approach, two parameters together define the strength of the path dependent effect:

- The penalty for divergence π defines how likely a second-period strategy is, given its distance from the first-period strategy
- The relative weight on path dependency w_p determines the probability of any second-period strategy as the combination of the optimal probability distribution (one if optimal in reference model, zero otherwise) and the pure path dependent distribution defined by π .

The results of the probabilistic path dependent model are shown in Figures 10 and 11. Figure 10 shows the sensitivity of optimal period 1 policy to the probability of high damages for three different levels of path dependency. These results assume a moderately stringent penalty for divergence π of 7. Note that the range of assumptions with a higher initial carbon tax is monotonically increasing in the weight on path dependency, unlike the cost approach in Figure 8. As the relative weight on path dependency decreases, the future decision is more highly constrained by the previous decision and is less likely to override that effect even if a strategy is optimal. The precise effect varies as a function of the probability of high damage, and within some ranges may be unchanged, but this is mainly an effect of the discretization of policy choices required to simulate the numerical models.

Figure 11 again displays the two-way sensitivity of varying both the probability of high damage and the relative weight on path dependency in period 2; the shading pattern indicates the optimal period 1 policy choice. Here again I assume a moderately stringent penalty for divergence π of 7. As in Figure 12, the

Figure 10 Optimal First-Period Strategy as the Probability of High Damage Is Varied, for Different Degrees of Path Dependency

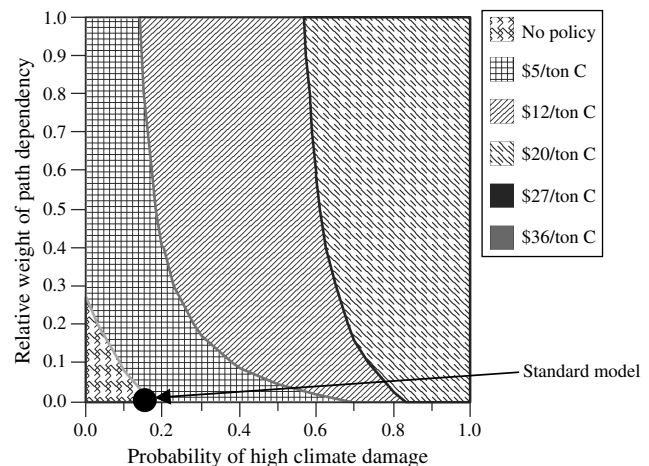


Note. Higher weight means more path dependency. $\pi = 7$.

block dot indicates the standard version of the model, with no path dependency and a probability of high damage of 0.185.

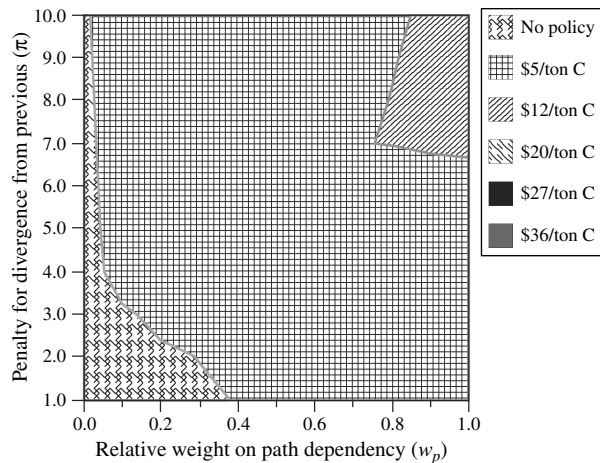
The effects of both parameters in the path-dependent model are shown in the two-way sensitivity analysis in Figure 12, assuming a probability of high climate-damage value of 0.185. This graph again shows that a small decrease in the weight on path dependency is enough to increase the optimal period 1 decision for almost any penalty for diver-

Figure 11 Sensitivity of Optimal Period 1 Decision on the Probability of High Damage and on the Weight of Path Dependency



Note. Black dot indicates the reference version of the model with no path dependency. $\pi = 7$.

Figure 12 Sensitivity of Optimal Period 1 Decision on the Weight on Path Dependency



Notes. Higher weight = more path dependence and on the penalty for diverging from the previous decision. Higher π means more severe path dependency. Probability of high damage valuation is 0.185.

gence, unless this is very weak ($\pi \leq 3$ has a nearly uniform distribution across all possible period 2 decisions—see Figure 3).

One effect of the reduced flexibility of future decisions from the path dependency is that the EVPI will be reduced. The EVPI measures the added gain if the uncertain parameters were known with certainty in period 1. However, with the decreased flexibility, the value of that knowledge is also decreased. We give the EVPI for several different strengths of the path dependent effect in Table 4.

Notice that although the qualitative behavior is the same under either a cost-based approach or a probabilistic approach to modeling path dependency, the probabilistic approach results in a higher range of carbon prices being optimal. The intuition for this is that under a cost-based approach, extreme reduc-

tions in period 2 may be expensive but are still feasible, whereas under the probabilistic approach, there is some probability, even if it is small, that stringent future reductions are not an available option. Higher carbon prices in period 1 increase the options available in the future, and thus have greater value under the probabilistic model of path dependency.

3.4. Sensitivity to Assumptions

To further test the significance of the path dependency effect on near-term policy, we explore the sensitivity of the results above to several variations. We present here variations to the probabilistic implementation of path dependency. One potential criticism to the probabilistic approach is that it may result in a bimodal distribution with a peak around the previous period's stringency and another peak around the optimal decision. To explore the sensitivity to the shape, we test two alternative versions of the model in which the probability distribution has been smoothed. The two versions smooth either by nearest-neighbor averaging or by nearest-two-neighbors averaging. An example probability distribution for a future decision is given in Table 5, along with the two smoothed versions. Note that the general behavior of higher initial carbon taxes still holds (Figure 13) and is in fact increased relative to the original path-dependent version.

A second question that may arise about the results from the previous section is how the results depend on the fact that the optimal policy in the nonpath-dependent version was to undertake no emissions reductions (\$0 carbon tax). To explore the effect of this assumption in the base model, we artificially lower the costs of emissions reductions everywhere proportionally to 50% of the original model values. This change will make more emissions reductions attractive even in the absence of path dependency, because the costs of climate damages have not

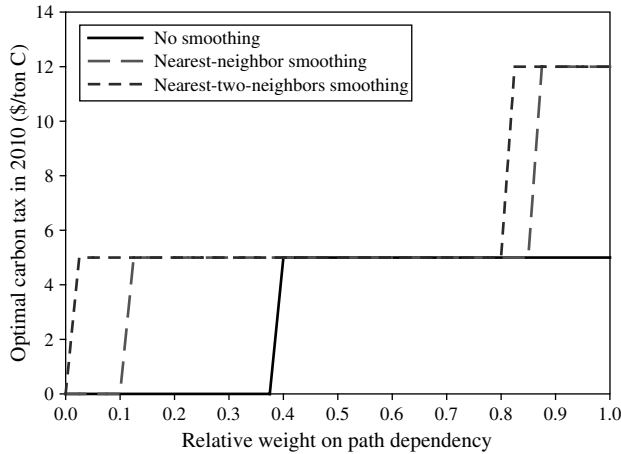
Table 4 Expected Value of Perfect Information as a Function of the Weight on Path Dependency in the Probabilistic Version

Weight on path dep.	EVPI (%)
0	0.19
0.2	0.16
0.4	0.14
0.6	0.12
0.8	0.11
1	0.09

Table 5 Example Probability Distribution of Period 2 Policy and Two Increasingly Smoothed Alternatives

Strategy (\$/ton C)	Original	Smoothed	Smoothed more
0	0.00	0.20	0.14
5	0.60	0.21	0.17
12	0.04	0.32	0.19
20	0.32	0.13	0.15
27	0.04	0.12	0.11
36	0.00	0.01	0.05

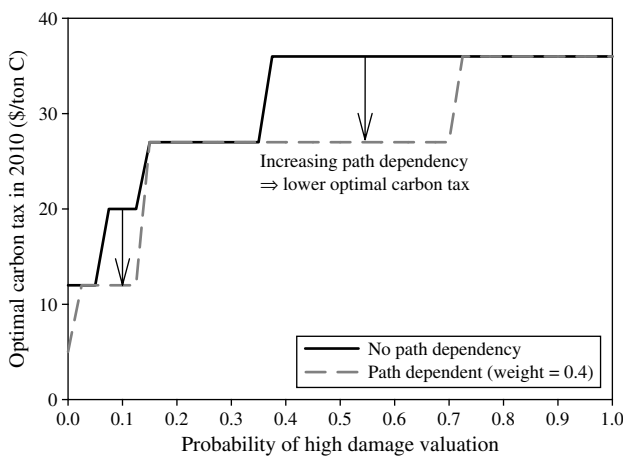
Figure 13 Sensitivity of Optimal Period 1 Decision on the Relative Weight on Path Dependency (Higher Weight = More Path Dependence)



Notes. Solid line shows the result from the original probabilistic version of the path dependent model, and the two dashed lines show the results where probabilities over future decisions are smoothed. The probability of high damage valuation is 0.185.

changed. We compare the optimal first-period decision in the model without path dependency to one with path dependency in Figure 14. Note that now, in the absence of path dependency, higher carbon taxes are optimal in the first period for almost any assumption about damages. In this model, adding path dependency has the opposite effect, resulting in a lower carbon tax in the first period. The general

Figure 14 Effect of Path Dependency on Optimal First-Period Strategy as the Probability of High Damage Is Varied, Low Abatement Cost Version



Note. $w_p = 0.4$ and $\pi = 7$

Table 6 Strategy Choices in Three-Period Model: Reduction Below Unconstrained Emission Growth Rate

Decision period	Strategy variable	Years	Strategy choices available (reduction below unconstrained rate) (%)
1	Policy 2010	2010–2029	0, 2.0, 2.5, 3, 3.5, 4.0
2	Policy 2030	2030–2049	0, 2, 4, 6, 8, 10
3	Policy 2050	2050–2100	0, 1, 2, 3, 4, 5

effect of path dependency is to induce greater hedging in the first period, with neither a very weak nor a very stringent policy, even if it is economically optimal.

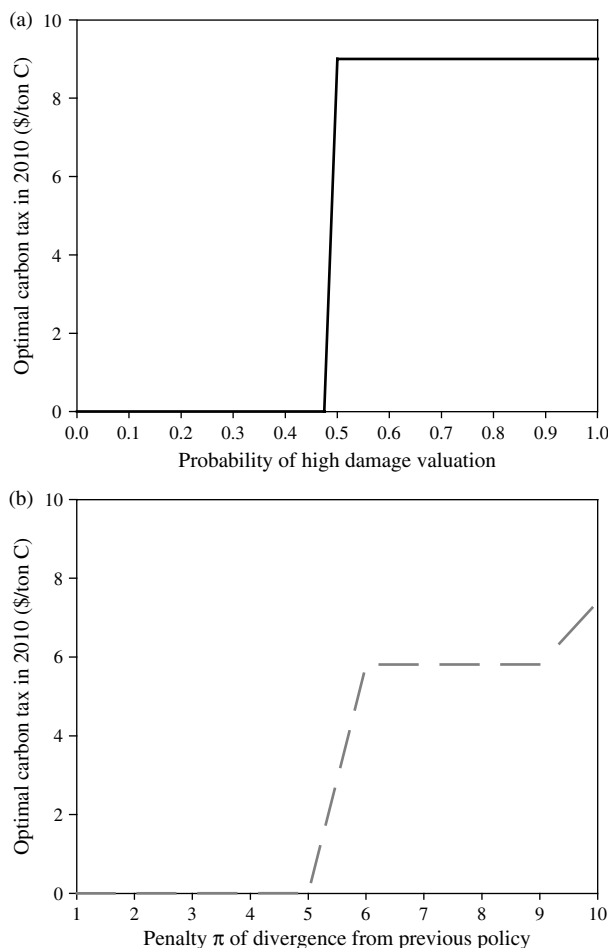
As a final sensitivity test, we explore the effect in a three-period model. The periods now consist of one from 2010 to 2030, another from 2030 to 2050, and a third from 2050 to 2100. The strategy space is as shown in Table 6, where the interpretation as before is as percentage point reductions in the emissions growth rate. In this version, I also assume that the uncertainty regarding the climate-damage valuation is revealed gradually. Between periods 1 and 2 (in 2030), a signal is received that makes low, medium, or high climate damage more likely, and the other values correspondingly less likely. Between periods 2 and 3 (in 2050), all remaining uncertainty is then resolved and the period 3 strategy is chosen under perfect certainty. The probability of receiving any signal in 2030 is the same as the original probability. If the signal is received for one damage level, the posterior probabilities of the other two levels decrease 50% from their prior, and the probability that the predicted level is revealed is increased by that same amount. The numerical values for the reference distribution are given in Table 7.

The optimal first-period policy in the new three-period decision model is shown in Figure 15(a) for the full range of possible priors for climate damage. Unlike the two-period model, the optimal policy in

Table 7 Prior And Posterior Probabilities for Damage Valuation Uncertainty in Three-Period Model

Posterior probability that damage is revealed to be:			Probability that signal in 2030 indicates that damage is		
Low	Med	High	Low	Med	High
0.5925	0.315	0.0925	0.185	0.63	0.185
0.0925	0.815	0.0925			
0.0925	0.315	0.5925			

Figure 15 Sensitivity Analysis of Optimal Period 1 Policy In Three-Period Model

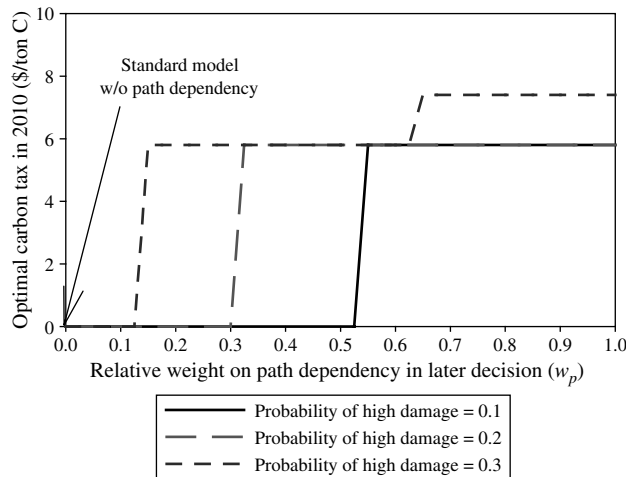


Notes. Graph (a) shows the sensitivity of the nonpath-dependent version to the probability of high damage valuation, and graph (b) shows the sensitivity of the path-dependent version to the penalty for divergence, with the weight of path dependency (w_p) fixed at 0.5 and the probability of high damage at 0.2.

2010 is never higher than \$9/ton, and is zero when the prior probability of high damage is less than 0.5. With the increased flexibility over later decades, near-term policy is even less justified.

I now add path dependency, as before, to this model to see whether consideration of this effect justifies a nonzero carbon price with lower probabilities of high damage. The decision nodes in periods 2 and 3 are converted to chance nodes, and the probability of each branch is calculated as in the previous model. In this three-period model with path dependence, the first period decision influences not only the policy

Figure 16 Sensitivity of First Period Decision to the Weight of Path Dependency, for Different Probabilities of High Damage, and Penalty $\pi = 7$



Note. The nonpath-dependent version is defined by $w_p = 0.0$

chosen in 2030, but also indirectly influences the policy chosen in 2050. For lower probabilities of high damage (e.g., 0.2), stronger path dependent effects are required to influence the period 1 policy. Figure 15(b) shows the optimal policy for different levels of the divergence penalty π . For values below $\pi = 5$, no carbon price is justified.

Fixing the divergence penalty π at 7, the relationship between optimal first-period policy and the relative weight on path dependency is given in Figure 16 for several prior probabilities of high climate damage. If path dependent constraints have even a moderate effect on future decisions, a carbon tax in the range of \$6–\$8 is justified for the next decade.

4. Discussion

Regardless of the implementation of the Kyoto Protocol in Europe and other participating developed nations, the current policy debate in both the United States and internationally is focused on what, if anything, should be done to restrict carbon dioxide and other greenhouse gases in the next decade or two. Some argue that mandatory regulations are not necessary yet, until we learn more and reduce the uncertainty about future climate effects. Formal analysis using models of decision under uncertainty have appeared to lend credence to this prescription.

The analysis presented above has modified a standard sequential decision model to include path dependency, a factor that has been ignored in formal climate policy models but is well recognized as an attribute of political systems. By representing path dependency either as a cost of departing from previous policies or as an influence over the probability of future policy shifts, I find that path dependence generally increases the optimal level of policy. In particular, where the standard model without path dependence has an optimal carbon price of zero, very slight path dependency using either method justifies a carbon price of at least on the order of \$5/ton C. Given the predominance of inertia and path dependency in political systems, the results of models that ignore path dependency are likely to be biased.

In developing decision-analytic models for policy applications, especially those occurring over long time horizons, ignoring the path dependent tendency in government decisions may well lead to near-term decisions that do not hedge sufficiently. Given the difficulty of enacting radical policy shifts, a near term strategy that assumes that a drastic change can be made later if needed is likely to be suboptimal. I have presented two alternative methods for including path dependency in decision-analytic models. The cost function approach is simpler and may be preferred for many applications, although there is little intuitive guidance for selecting the magnitude of the cost penalty. The probabilistic approach is a reasonable alternative, especially when the time horizons are so long that the future decision makers are not the same as the current generation, and the relative influence of optimality versus status quo on decisions may be a conceptually easier quantity on which to elicit expert judgments.

The general result of including path dependency in a decision model is much like the quasi-option value of Arrow and Fisher. The reason is that path dependency acts like another type of irreversibility. Just as the earth's climate system and capital stocks in the economy have irreversibilities, so too do national and international political institutions. Considering this additional irreversibility will cause a rational forward-looking agent to increase the hedging strategy against the possibility of required future policy

action. Note that this effect can work in either direction; if the optimal level of effort without path dependence is very large, path dependence will indicate that initial policy should be slightly more moderate.

The conclusion of this study for the climate policy modeling community is that applications of sequential decision models over very long time horizons should consider path dependencies in the political systems modeled. Otherwise, if policies at each time point can be reconsidered without regard to past decisions, we may place an unrealistic expectation on future generations and eliminate future options by not laying the groundwork with minimal policies today. Ignoring path dependencies risks giving qualitatively biased advice to policymakers as to whether it is yet time to begin mandatory emissions regulations.

Acknowledgments

The author gratefully acknowledges helpful comments from Gary Yohe, Henry Jacoby, Christopher Robert, the students of UNC Public Policy Department, and two anonymous reviewers, and the associate editor of *Decision Analysis*. The work also benefited from the able assistance of Jennifer Holak. This research was supported by the U.S. Department of Energy (DE-FG02-02ER63468, DE-FG04-04ER63929), the U.S. Environmental Protection Agency, the U.S. National Science Foundation, the U.S. National Aeronautics and Space Administration, the U.S. National Oceanographic and Atmospheric Administration; and the Industry and Foundation Sponsors of the MIT Joint Program on the Science and Policy of Global Change.

References

- Altman, M. 2000. A behavioral model of path dependency: The economics of profitable inefficiency and market failure. *J. Socio-Econom.* **29** 127–145.
- Arrow, K. J., A. C. Fisher. 1974. Environmental preservation, uncertainty, and irreversibility. *Quart. J. Econom.* **88** 312–319.
- Babiker, M. H., H. D. Jacoby, J. M. Reilly, D. M. Reiner. 2002. The evolution of a climate regime: Kyoto to Marrakech. *Environ. Sci. Policy* **5**(3) 195–206.
- Bertsimas, D., L. Kogan, A. W. Lo. 2001. Hedging derivative securities and incomplete markets: An epsilon-arbitrage approach. *Oper. Res.* **49**(3) 372–397.
- Bingaman, J. 2007. Low carbon economy act of 2007. Senate Bill 1766. <http://bingaman.senate.gov/news/stlvrecord.cfm?id=284811>.
- Bush, G. W. 2001. Text of a letter from the President to Senators Hagel, Helms, Craig, and Roberts. Available at <http://www.whitehouse.gov/news/releases/2001/03/20010314.html>.
- Bush, G. W. 2002. *Global Climate Change Policy Book*. Available at <http://www.whitehouse.gov/news/releases/2002/02/climatechange.html>.
- Dixit, A. 1991. Analytical approximations of hysteresis. *Rev. Econom. Stud.* **58**(1) 141–151.

- Economist, The*. 2006. Climate change: Hot under the collar. (April 1) 33–34. Available at http://www.economist.com/World/europe/displayStory.cfm?story_id=6746635.
- Ellerman, A. D., P. L. Joskow, R. Schmalensee, J.-P. Montero, E. M. Bailey. 2000. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge University Press, Cambridge, UK, 362.
- Forest, C. E., P. H. Stone, A. P. Sokolov, M. R. Allen, M. D. Webster. 2002. Quantifying uncertainties in climate system properties using recent climate observations. *Science* **295**(5552) 113–117.
- Gerlagh, R., M. A. Keyzer. 2004. Path-dependence in a Ramsey model with resource amenities and limited regeneration. *J. Econom. Dynam. Control* **28** 1159–1184.
- Hacker, J. 1998. The historical logic of national health insurance: Structure and sequence in the development of British, Canadian, and U.S. medical policy. *Stud. Amer. Political Development* **12**(Spring) 57–130.
- Ha-Duong, M. 1998. Quasi-option value and climate policy choices. *Energy Econom.* **20** 599–620.
- Hammit, J. K., R. A. Lempert, 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. *Amer. Econom. Rev.* **64**(6) 1006–1012.
- Jacoby, H. D., R. Eckaus, A. D. Ellermann, R. Prinn, D. Reiner, Z. Yang. 1997. CO₂ emissions limits: Economic adjustments and the distribution of burdens. *Energy J.* **18** 31–58.
- Kaplow, L., S. Shavell. 2002. On the superiority of corrective taxes to quantity regulation. *Amer. Law Econom. Rev.* **4**(1) 1–17.
- Keefer, D., S. E. Bodily. 1983. Three-point approximations for continuous random variables. *Management Sci.* **40** 760–773.
- Kolstad, C. D. 1996. Learning and stock effects in environmental regulation: The case of greenhouse gas emissions. *J. Environ. Econom. Management* **31** 1–18.
- Levi, M. 1997. A model, a method, and a map: Rational choice in comparative and historical analysis. M. I. Lichbach, A. S. Zuckerman, eds. *Comparative Politics: Rationality, Culture, and Structure*. Cambridge University Press, Cambridge, 19–41.
- Liebermann, J. 2007. America's climate security act of 2007. Senate Bill 2191. Available at <http://www.govtrack.us/congress/bill.xpd?bill=s110-2191>.
- Lipset, S. M., S. Rokkan. 1967. Cleavage structures, party systems and voter alignments: An introduction. S. M. Lipset, S. Rokkan, eds. *Party Systems and Voter Alignments*. Free Press, New York, 1–64.
- Manne, A. S., R. G. Richels. 1992. *Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emission Limits*. MIT Press, Cambridge, MA.
- Manne, A. S., R. G. Richels. 1995. The greenhouse debate: Economic efficiency, burden sharing and hedging strategies. *Energy J.* **16**(4) 1–37.
- McFarland, J. R., J. M. Reilly, H. J. Herzog. 2004. Representing energy technologies in top-down economic models using bottom-up information. *Energy Econom.* **26**(4) 685–707.
- National Commission on Energy Policy (NCEP). 2004. Ending the energy stalemate: A bipartisan strategy to address America's energy challenges. National Commission on Energy Policy, Washington, DC, available at <http://www.energycommission.org/materials/>.
- Nordhaus, W. D. 1994a. *Managing the Global Commons: The Economics of Climate Change*. MIT Press, Cambridge, MA.
- Nordhaus, W. D. 1994b. Expert opinion on climatic change. *Amer. Scientist* **82**(January) 45–51.
- Nordhaus, W. D., D. Popp. 1997. What is the value of scientific knowledge? An application to global warming using the PRICE model. *Energy J.* **18**(1) 1–45.
- Paltsev, S., J. M. Reilly, H. D. Jacoby, R. S. Eckaus, J. McFarland, M. Sarofim, M. Asadoorian, M. Babiker. 2005. The MIT emissions prediction and policy analysis (EPPA) model: Version 4. MIT Joint Program on the Science and Policy of Global Change, Report 125. Available at <http://web.mit.edu/globalchange/www/reports.html#pubs>.
- Pierson, P. 2000. Increasing returns, path dependence, and the study of politics. *Amer. Political Sci. Rev.* **94**(2) 251.
- Pizer, W. A. 1999. The optimal choice of climate change policy in the presence of uncertainty. *Resource Energy Econom.* **21**(3, 4) 255–287.
- Prinn, R., H. Jacoby, A. Sokolov, C. Wang, X. Xiao, Z. Yang, R. Eckhaus, P. Stone, D. Ellerman, J. Melillo, J. Fitzmaurice, D. Kicklighter, G. Holian, Y. Liu. 1999. Integrated global system model for climate policy assessment: Feedbacks and sensitivity studies. *Climatic Change* **41**(3/4) 469–546.
- Reilly, J. M., S. Paltsev. 2005. An analysis of the European emission trading scheme. MIT Joint Program on the Science and Policy of Global Change. Report 127, Cambridge, MA.
- Reilly, J. M., M. Mayer, J. Harnisch. 2002. The Kyoto protocol and non-CO₂ greenhouse gases and carbon sinks. *Environ. Model. Assessment* **7**(4) 217–229.
- Revkin, A. 2005. Climate debate: Incremental gains. *The New York Times* (December 11).
- Roughgarden, T., S. H. Schneider. 1999. Climate change policy: Quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy* **27** 415–429.
- Sewell, W. H. 1996. Three temporalities: Toward an eventful sociology. T. J. McDonald, ed. *The Historic Turn in the Human Sciences*. University of Michigan Press, Ann Arbor, 245–263.
- Sokolov, A. P., P. H. Stone. 1998. A flexible climate model for use in integrated assessments. *Climate Dynam.* **14** 291–303.
- Tatang, M. A., W. Pan, R. G. Prinn, G. J. McRae. 1997. An efficient method for parametric uncertainty analysis of numerical geophysical models. *J. Geophysical Res.* **102**(D18) 21,925–21,932.
- Ulph, A., D. Ulph. 1997. Global warming, irreversibility and learning. *Econom. J.* **107**(442) 636–650.
- Valverde, L. J., Jr., H. D. Jacoby, G. Kaufman. 1999. Sequential climate decisions under uncertainty: An integrated framework. *J. Environ. Model. Assessment* **4** 87–101.
- Webster, M. D. 2002. The curious role of learning: Should we wait for more data? *Energy J.* **23**(2) 97–119.
- Webster, M. D., A. P. Sokolov. 2000. A methodology for quantifying uncertainty in climate projections. *Climatic Change* **46**(4) 417–446.
- Webster, M. D., M. Babiker, M. Mayer, J. M. Reilly, J. Harnisch, M. C. Sarofim, C. Wang. 2002. Uncertainty in emissions projections for climate models. *Atmospheric Environment* **36**(22) 3659–3670.
- Webster, M. D., C. Forest, J. Reilly, M. Babiker, D. Kicklighter, M. Mayer, R. Prinn, M. Sarofim, A. Sokolov, P. Stone, C. Wang. 2003. Uncertainty analysis of climate change and policy response. *Climatic Change* **61**(3) 295–320.
- Weimer, D. L., A. R. Vining. 2005. *Policy Analysis: Concepts and Practice*, 4th ed. Pearson Prentice Hall, Upper Saddle River, NJ.
- Yohe, G., N. Andronova, M. Schlesinger. 2004. To hedge or not against an uncertain climate future. *Science* **306** 416–417.
- Zhao, J., C. L. Kling. 2003. Policy persistence in environmental regulation. *Resource and Energy Econom.* **25** 255–268.