
An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty

Mort Webster · Nidhi Santen · Panos Parpas

March 15, 2012

Abstract Analyses of global climate policy as a sequential decision under uncertainty have been severely restricted by dimensionality and computational burdens. Therefore, they have limited the number of decision stages, discrete actions, or number and type of uncertainties considered. In particular, two common simplifications are the use of two-stage models to approximate a multi-stage problem and exogenous formulations for inherently endogenous or decision-dependent uncertainties (in which the shock at time $t + 1$ depends on the decision made at time t). In this paper, we present a stochastic dynamic programming formulation of the Dynamic Integrated Model of Climate and the Economy (DICE), and the application of approximate dynamic programming techniques to numerically solve for the optimal policy under uncertain and decision-dependent technological change in a multi-stage setting. We compare numerical results using two alternative value function approximation approaches, one parametric and one non-parametric. We show that increasing the variance of a symmetric mean-preserving uncertainty in abatement costs leads

Mort Webster
Engineering Systems Division
Massachusetts Institute of Technology
77 Massachusetts Avenue Cambridge, MA 02139
E-mail: mort@mit.edu

Nidhi Santen
Engineering Systems Division
Massachusetts Institute of Technology
77 Massachusetts Avenue Cambridge, MA 02139
E-mail: nrsanten@mit.edu

Panos Parpas
Department of Computing
Imperial College London
180 Queen's Gate
London, SW7 2AZ
United Kingdom
E-mail: p.parpas@imperial.ac.uk

to higher optimal first-stage emission controls, but the effect is negligible when the uncertainty is exogenous. In contrast, the impact of decision-dependent cost uncertainty, a crude approximation of technology R&D, on optimal control is much larger, leading to higher control rates (lower emissions). Further, we demonstrate that the magnitude of this effect grows with the number of decision stages represented, suggesting that for decision-dependent phenomena, the conventional two-stage approximation will lead to an underestimate of the effect of uncertainty.

Keywords Climate policy analysis · approximate dynamic programming · decision dependent uncertainty · stochastic dynamic programming · endogenous uncertainty

1 Introduction

Responding to the threat of global climate change is one of the most difficult risk management problems that society faces. An optimal path of greenhouse gas emissions reductions in principle should be the path that balances the costs of emissions reductions, or abatement, with the climate-related damages from emissions. However, both the costs of emissions reductions and the damages from climate change are uncertain, and neither will be known with certainty for a long time. Nevertheless, information about the uncertainties will be revealed gradually, and policies will be continually responding to new information and other changing conditions.

Models that represent the complete causal chain from economic activity to ultimate physical impacts of climate change are referred to as “integrated assessment models (IAMs).” They simulate both the economic and the biogeophysical systems and their interactions in a single model. The majority of analyses with integrated assessment models are deterministic, and are focused on understanding and improving representations of the integrated system. There has been some work applying probabilistic uncertainty analysis to IAMs, usually in the form of Monte Carlo simulation, e.g., [16, 26, 27, 32, 33]. Studies that have explicitly modeled sequential decision under uncertainty have represented the problem in a highly simplified and stylized manner, often as a two-stage problem with a small number of discrete actions and uncertainties (e.g., [2, 10, 13, 29–31, 37]).

An appropriate framing of this problem is as a dynamic stochastic optimization model. This general class of problems can be formulated and solved with either stochastic programming with recourse or dynamic programming methods. There are several special challenges to the climate problem that make it difficult to solve with existing numerical methods. First, the long time-lags in the earth system make it necessary to simulate at a minimum a few centuries ahead. Policy decisions can be revised at any time, making this a decision problem with many stages. The dimensionality of the problem is increased further by the number of uncertainties inherent in projecting global economic and technological change over several centuries and in projecting the response

of the earth’s climate system to greenhouse gas emissions. The action space (emissions reductions) and state space (all variables required to describe the evolution of the system over time) are continuous variables that may need to be discretized. The dimensionality of this problem, even in a highly simplified form, is extremely large.

One important complication associated with current integrated assessment models is that arguments for near-term emissions reductions are motivated less by the value of the emissions they avoid in this long-term problem and more by the possibility that policies today will encourage technological change which will lower future abatement costs. To explore this argument in a rigorous framework requires modeling endogenous or decision-dependent uncertainties, since the decision to abate today will change the probability distribution of next period’s abatement costs. Modeling decision-dependencies of this type poses a unique challenge to conventional stochastic programming methods, which typically use exogenous scenario trees. Existing stochastic programming methods that attempt to model decision-dependent uncertainties cannot be applied to this model because they apply only under very specific circumstances. For example, Goel and Grossman [9] present a framework in which decisions affect the time in which the uncertainties will be resolved. Baker and Solak [2] introduce a stochastic programming version of an IAM with endogenous decision-dependent probabilities, but one that uses a customized deterministic mapping function to assign outcomes to decisions. Decisions in climate policy analysis influence the probability of different outcomes, and the need to have a flexible way to capture endogenous uncertainties means that Stochastic Dynamic Programming (SDP) is an appropriate framework for climate policy analysis. Unfortunately, classical SDP algorithms (e.g., value iteration, policy iteration [4]), suffer from the curse of dimensionality; i.e., the complexity of the problem grows exponentially with the number of states.

There have been a few studies that have formally framed the climate decision problem under uncertainty as a multi-stage stochastic dynamic program, using a variety of approaches to overcome the dimensionality challenge. Gerst et al. [8] use discrete sampling via experimental design and a very large number of iterations to learn about the solution space, which can be computationally expensive. Kelly and Kolstad [12] and Leach [14] approximate the value function associated with the Bellman equation using neural networks to estimate a functional form with 16 terms, but use discrete gridded samples in state-space to iteratively improve the approximation. Crost and Traeger [6] and Lemoine and Traeger [15] statistically estimate relationships between state variables offline in order to reduce the dimensions of the state vector, and then use conventional backward induction on the reduced state-space. All of these approaches rely on discretizing a (possibly reduced) state-space into intervals, and therefore require difficult tradeoffs between resolution/accuracy and computation time.

Here we present an alternative efficient solution method for multi-stage, multi-dimensional stochastic dynamic programs, based on Approximate Dynamic Programming (ADP)[4,25]. In this approach, we approximate the value

function with a continuous function, which avoids the resolution and computational issues of discretized approaches. A key challenge associated with the successful application of the ADP methodology is the specification of the set of basis functions used to construct an approximate value function. The solution obtained via ADP methods is known to be sensitive to the choice of basis functions that are used to build the value function. If the true value function is not spanned by this basis then the ADP algorithm will converge to the wrong solution. This false convergence is difficult to detect in practice.

We address this issue using two alternative approaches for value function approximations, one parametric, using global regression, and one non-parametric, using a mesh-free moving least squares approach. The parametric method is in principle faster but may exhibit the false convergence issue discussed above. The non-parametric method may be slower but can be used to detect errors in the choice of basis functions. We develop and test our algorithm using a stochastic dynamic programming version of the Dynamic Integrated model of Climate and the Economy (DICE) [20]. We demonstrate that for this application, ADP has several advantages over alternative solution methods including the ability to model decision-dependent uncertainties, manage a high-dimensional state space over a multi-stage stochastic decision problem, and converge in a fraction of the computational time. Using numerical results, we show that an increase in uncertainty in future abatement costs leads to a slight increase in the optimal level of near-term emissions reductions when the uncertainty is exogenous. However, once a probabilistic decision-dependent effect is included, the optimal near-term emissions reductions are greater than the expected value case. In this latter context, we demonstrate that the effect of decision-dependent uncertainty increase with the number of decision stages represented, suggesting a potential bias in two-stage approximations.

We describe the DICE model, the formulation of the stochastic version, and the algorithms for solution using ADP in Section 2. Section 3 validates the new algorithms. In Section 4, we present the results of simulations of the base model using both parametric and non-parametric value function approximations, as well as the results of the decision-dependent variation. Section 5 gives a concluding discussion and suggests directions for future research.

2 Methods

Integrated assessment models [17, 34] are a general class of models that couple economic growth equations with differential equations that describe the transient evolution of the biogeophysical earth system. IAMs fall into two broad subgroups: policy evaluation models, which simulate exogenous emissions policies, and policy optimization models, which are optimal control models. The model described in this study falls into the latter category. We illustrate our computational solution algorithm on one such model, but the techniques are broadly adaptable to other models in the same class as well as other intertemporal optimization models.

2.1 The DICE Model

The effect of learning on optimal policy choice is calculated using a stochastic version of the DICE-99 model [20]¹. The DICE-99 model is a Ramsey growth model augmented with equations for CO₂ emissions as a function of economic production, the carbon-cycle, radiation, heat balance, abatement cost and climate damage cost functions. The model solves for the optimal path over time of the savings/consumption decision, and also the emissions abatement decision that balances the cost of emissions abatement against damages from increased temperatures. Specifically, DICE is a deterministic, constrained non-linear program which chooses investment $I(t)$ and abatement $\mu(t)$ in order to maximize the sum of discounted utility:

$$\max_{I(t), \mu(t)} \sum_{t=0}^T U(c(t), L(t))(1 + \rho(t))^{-1}, \quad (1)$$

where $U(\cdot, \cdot)$ is the utility function, $c(t)$ is the per capita consumption, $L(t)$ is the population, and $\rho(t)$ is the social rate of time preference. The maximization is subject to a set of constraints, including the production function for the economy, the relationship between economic output and emissions, relationships for concentrations, radiative forcing, temperature change, and the reduction in output from both abatement costs and damage costs. The full set of equations for the model is given in the Appendix below and more details are in [20]. The time horizon of the model is 350 years in 10 year steps.

2.2 Formulation of the Decision under Uncertainty Problem

Parameters in the DICE model are uncertain, as clearly we do not have perfect information about future economic growth and technological change or a complete understanding of the earth's climate system. But the uncertainty in some parameters are more important than others in terms of their effect on optimal abatement decisions. Nordhaus and Popp [21] performed uncertainty analysis of the DICE model and concluded that the most critical parameters are those that determine the costs of emissions reductions and those that determine the ultimate damages from temperature changes (as opposed to uncertainties in baseline projections). Decisions under uncertainty in climate damages have been more fully examined by others [6, 10, 13, 29, 37]. However, Kelly and Kolstad [12] and Webster et al. [31] have shown that it may take a very long time before the uncertainty in damages is reduced. In contrast, some of the uncertainty in the costs of emissions reductions may be reduced sooner. The complication with cost uncertainty is that the time in which new information is obtained will depend on the level of abatement attempted. In this analysis, we focus on the uncertainty in the cost of abatement. The problem becomes

¹ Newer versions of DICE exist [19], however we use DICE-99 for its relative simplicity and because the subsequent updates do not change the qualitative points being made here.




Fig. 1 Schematic of Sequential Abatement Decision under Uncertainty

one of choosing a level of abatement in each decision stage under uncertainty in abatement costs, after which information is received about costs that may shift the expected future costs, which are still uncertain. In the next stage, abatement levels are chosen again after observing the realized costs in the previous period. The decision problem under uncertainty is illustrated in Figure 1, using a skeleton of a decision tree, assuming a finite set of N decision stages. Mathematically, the stochastic problem is that in equation (2), where μ_t is the abatement level in stage t , θ_t is the cost shock in stage t , and R_t is the discounted utility in stage t .

$$\max_{\mu_1} \left\{ R_1 + \max_{\mu_2} E_{\theta_1} [R_2 + \dots] \right\}. \quad (2)$$

We implement and solve the stochastic version of the model using the framework of stochastic dynamic programming. Dynamic programming uses the Bellman equation [3] to decompose the problem in (2) into the relatively simpler set of conditions that must hold for all decision stages t :

$$V_t = \max_{\mu_t} [R_t + E \{V_{t+1}(\mu_t, \theta_t)\}]. \quad (3)$$

The expression in (3) captures the essential nature of the stochastic optimization problem; just as the deterministic DICE finds the intertemporal balance between the costs of reducing emissions now and the future benefits of avoided climate damage, the stochastic problem is to find the optimal balance between near-term costs and expected future costs. As an illustration, Figures 2 and 3 show the near-term costs and expected future costs, respectively, that are balanced in the first of the two numerical experiments presented below.

In the deterministic version of DICE, abatement cost as a percentage of output (GDP) is a function of the abatement decision variable,

$$AC = 1 - b_1 \mu^{b_2}.$$

Where c_1 and c_2 are calibrated as in [20]. b_1 starts at a value of 0.03 and grows over time. The growth rate of the cost coefficient declines over time as

$$g_b(t) = -0.08e^{-0.08t}.$$

The cost coefficient grows as,

$$b_1(t) = \frac{b_1(t-1)}{(1 - g_b(t))}.$$




Fig. 2 Current stage utility as a function of control rate in first decision stage when abatement cost is uncertain.




Fig. 3 Expected future utility as a function of control rate in first decision stage when abatement cost is uncertain.

We represent uncertainty in future abatement costs as an i.i.d. multiplicative shock in each period to the reference growth rate of costs $g_c(t)$. Because the shock is applied to the growth rate of costs, the uncertainty in the level of costs have memory; i.e., the sample paths of abatement costs over time are not mean-reverting. In the results presented below, the reference distribution for the cost growth rate shock is assumed to be Normal with a mean of 1.0 and a standard deviation of 0.4. The uncertainty in abatement cost is based on a detailed uncertainty analysis of a higher resolution economic model of emissions [32]. Below, we also present results from a sensitivity analysis of the optimal control to the standard deviation of the cost shock.

For ease of exposition, we have made a few simplifying assumptions in the stochastic decision model. First, we fix the savings rate in capital stock in the economy $I(t)/Y(t)$ to the optimal trajectory from the deterministic model. This is because the optimal rate of investment is largely unresponsive to changes in abatement decisions or in abatement cost assumptions. The DICE model is defined in 10-year steps over a 350-year time horizon (35 model periods). Rather than define each decade as a decision stage, each decision stage in the stochastic model consists of 5 DICE model periods (50 year steps), for a total of seven decision stages ($N = 7$). Multi-decade decision stages, as opposed to decadal, make the communication of results easier and more importantly better characterize the stochastic process being modeled. The lifetime of many of the large capital investments that are affected by the abatement decision, such as a coal-fired power plant, typically have lifetimes of 30 – 50 years. Similarly, information about technological breakthroughs that drive the changes in the abatement costs may not occur every 10 years, nor do large policy shifts. The seven-stage model presented here approximates the time scale of the problem, while having significantly higher resolution than the two-stage model approaches that are most common in the literature.

Our primary objective here is to explore the impacts of decision-dependence. In the example here of abatement costs, much of the uncertainty in costs of reducing future emissions is due to uncertainty in the success of technological change. One conception of technological change is called learning-by-doing [1, 36,35], which views the rate of progress (e.g., in cost reductions) as a function of cumulative installed capacity or use. This view of technological change is clearly decision-dependent. For example, significant efforts to reduce greenhouse gas emissions would lead to greater capacities of low-carbon emitting technologies, such as wind or carbon-capture and storage, which in turn might drive down the costs of those technologies relative to a world without the emission reduction efforts. An alternative representation of technological change is learning-by-searching, such as explicit economic models of R&D [23, 11]. In this latter process, the feedbacks between emissions policy and cost reductions are indirect and more controversial, but plausible enough to be a common motivation for those seeking such policies. Such a causal chain would link emission reduction policies, and their corresponding relative price effects, to an incentive to direct greater R&D expenditures towards lower-emitting technologies, and the greater expenditures are assumed to have more success on average at lowering costs. This feedback is more controversial, particularly the final causal step in the preceding description. Our purpose in this study is not to develop a detailed representation of technological change, but rather to focus on the process of decision-dependence which exists in some fraction of the full set of complex feedbacks that emissions reduction efforts would stimulate. Motivated by the type of decision-dependence suggested by some models of technological change, we developed an alternative version of the stochastic model using a stylized version of the general feedback mechanism. Specifically, we assume that abatement μ in one decision stage lowers the mean of the cost distribution in the next stage as

$$\begin{aligned}\tilde{c}_1(t) &= \tilde{c}_1(t-1)(1 - \alpha\mu(t-1)) & t > 1 \\ \tilde{c}_1(t) &= \bar{c}_1(t) & t = 1\end{aligned}$$

where \tilde{c}_1 is the mean of the cost distribution in the decision-dependent version, \bar{c}_1 is the mean of the cost distribution in the reference model, and $\alpha > 0$ is a scaling constant that alters the magnitude of the decision-dependent effect. We assume memory in this stochastic process; i.e., additional controls implemented in the next period can further lower the current mean of abatement costs. This representation will necessarily lead to probability distributions of abatement costs under the optimal policy that have a lower mean than the reference version of the model, which will in turn lead to higher emission control rates simply as a function of the lower costs. In order to control for this effect, the comparisons below between the decision-dependent and the exogenous models use a modified set of cost distributions in the exogenous version. Specifically, the decision-dependent version is simulated for several thousand iterations beyond convergence, and the abatement cost samples that result

under the optimal decision-dependent policy are used to construct new distributions. These distributions of cost uncertainty, which vary over decision stages, are then sampled i.i.d in the version denoted as 'exogenous'.

2.3 Approximate Dynamic Programming Implementation

The finite-horizon stochastic dynamic programming problem formulated above is traditionally solved as a Markov Decision Problem [4], using a backward induction algorithm. The algorithm iterates over the state, action, and uncertainty spaces for each decision stage to calculate the exact value function and corresponding policy function in each decision stage. Because the action and state spaces are all continuous, this would require discretization for each variable. For the DICE model, there are seven state variables that must be known; the capital stock ($K(t)$), three carbon concentration variables for a three-box model of the carbon cycle ($MAT(t)$, $MU(t)$, $ML(t)$), two temperature variables for a two-box energy-balance model ($TE(t)$, $TL(t)$), and the evolving abatement cost coefficient ($c_1(t)$). All of these variables require knowledge of the previous value to calculate the next value (see [20]).

In addition to the state variables, conventional dynamic programming would also iterate over discretized values of the action $\mu(t)$ and the cost growth shock $\theta(t)$, resulting in at least a 9-dimensional problem in each of 7 decision stages. This is an extremely large problem even if the discrete intervals are at an unsatisfyingly coarse resolution.

Instead of traditional backward induction, we have developed an approximate dynamic programming (ADP) algorithm for solving this problem, shown in Algorithm 1 (reference to only one of the two value function approximations is made). ADP is a family of methods (e.g., [5, 25]) that approximates the value function in each stage by adaptively sampling the state space to focus on higher expected value states until the value function converges. One critical advantage of forward sampling is that this enables a straightforward representation of decision-dependency. Two critical design choices in any efficient ADP algorithm are 1) the sampling strategy, and 2) the value function approximation.

Our solution algorithm consists of two phases. In phase I, the bootstrap phase, we use Latin Hypercube Sampling [18] to explore both the action space over all stages and the cost shock space. These sample paths are simulated forward, and the resulting Bellman values for the sample states and actions are saved for each decision stage. The full set of these samples of the value function are used to produce the first estimate of the value function approximation for each decision stage, using either of the two methods described below.

In phase II, we randomly sample the cost shock in each period to obtain a sample path, and choose the optimal action in each stage using the current value function approximations for the value of the next state, and the simulated DICE equations to obtain the current reward. The overall sampling approach

Algorithm 1: DICE Approximate Dynamic Programming Algorithm

Input: Decision stages N , bootstrap iterations bs , possible controls μ , uncertainty variable $\theta \sim N(1, \sigma)$, system state $s_0 \in S$ at time t_0 , system state transition equations $F(\mu, \theta)$, convergence criterion, $\bar{\epsilon}$

Phase I Initialization-Bootstrap: While $i \leq bs$,

1. Forward Pass

Loop over t from **1** to N , Latin Hypercube Sampling from θ and μ and set current reward as:

$$R_t(s_i) = U(c_t, L_t)(1 + \rho_t)^{-1}.$$

2. Backward Pass

Loop over t from N to **1**, setting the Bellman Value as:

$$v_t(s_i) = (R_t(s_i) + v_{t+1}(y_i|s_i))$$

where y_i is the sampled next system state resulting from μ_t and θ_t , and v_N is a pre-defined terminal value.

3. Construct First Estimate of Value Function: When $i = bs$, use OLS to set:

$$\hat{v}_t(s) = \Phi(s)r_0,$$

where Φ is a row vector of basis functions and r_0 is a column vector of coefficients that solves:

$$\min_{r_0} \sum_{s_i} (\hat{v}_t(s_i) - \Phi(s_i)r_0)^2.$$

for all sample states s_i .

Phase II Main Loop-Optimization: While $i > bs$,

1. Forward Pass

Loop over t from **1** to N , sampling θ randomly and sampling controls μ that achieve:

$$\max_{\mu} [R_t(s_i) + E\{v_{t+1}(y_i|s_i)\}]$$

where

$$E\{v_{t+1}(y_i|s_i)\} = \hat{v}_{t+1}(\mu_t, \theta_t).$$

Set current reward, $R_t(s_i)$, as in Phase I.

2. Backward Pass

Loop over t from N to **1**, setting the new Bellman Value as:

$$v_t(s_i) = (R_t(s_i) + \hat{v}_{t+1}(y_i|s_i))$$

where y_i is the sampled next system state.

Update r_i using a Bellman Error routine:

$$r_{i+1} = r_i - \gamma_i \epsilon_i \nabla r_i$$

where γ_i is a predefined smoothing parameter and

$$\epsilon_i = v_t(s_i) - \hat{v}_t(s_i).$$

Exit when:

$$\bar{\epsilon} = |\bar{v}_{1,i} - \bar{v}_{1,i-1}|$$

where $\bar{\epsilon}$ represents the change in the moving average of the total Bellman value in the initial stage.

Output: Optimal first-stage control, μ_1^* , value function approximations, $v_t^*(s)$

is an efficient (stratified) pure explore strategy in Phase I and a pure exploit strategy in Phase II.

In this study, we compare two alternative approaches to value function approximation, one parametric and one non-parametric. Both approaches approximate the expected value of being in any state as a reduced-form function of key features. Because of the forward sampling, not all state variables required for backward induction are needed as the key features or basis functions [5]. For this application, the fundamental structure is one of balancing near-term costs of abatement (reducing the size of the economy) against long-term costs from climate change. In terms of the state variables described above, the key features needed to approximate the value function are the capital stock $K(t)$ and the global surface temperature change $TE(t)$. The parametric approach employed is an iterative least squares regression method [4], approximating the value function as a nonlinear function of capital stock and temperature change. That is, the approximation of the value function is

$$\hat{v}_t(s) = \Phi(s)r \quad (4)$$

where Φ is a row vector of basis functions and r is a column vector of coefficients that solves,

$$\min_r \sum_{s_i} (\hat{v}_t(s_i) - \Phi(s_i)r)^2.$$

for all sample states s_i . Given an initial estimate of the coefficient vector r from the bootstrap phase, we iteratively improve the estimate using a Bellman error approach [4].

We compare an iterative regression approach with a non-parametric alternative. In this second approach, we apply moving least squares (MLS) [7] to interpolate the value function at a given state within a neighborhood. Meshfree methods such as MLS have been applied to other problems requiring interpolation in high dimensional space such as scattered data modeling, the solution of partial differential equations, medical imaging, and finance [7]. In the context of stochastic optimization, MLS was applied in [22] in an iterative algorithm that solves for the stochastic maximum principle. Here we apply the method in the context of the dynamic programming principle.

The approximate value of a state s is:

$$\hat{v}_t(s) = \bar{\Phi}(s)\bar{r}(s) \quad (5)$$

The difference between equations (4) and (5) is that the coefficient vector \bar{r} depends on the state s . Note that $\bar{r}(s)$ is obtained by solving

$$\min_{\bar{r}} \sum_{s_i} (\hat{v}_t(s_i) - \bar{\Phi}(s_i)\bar{r}(s_i))^2.$$

for all sample states s_i within some neighborhood of the state s . This requires solving many regressions, one for each point to be interpolated, as compared with the parametric approach. However, these regressions are generally for a

small number of samples in the immediate neighborhood, whereas the parametric approach is global and uses all samples which can grow to a large number. Thus, the tradeoff is between solving many small regressions (MLS) versus fewer larger regressions (parametric). Furthermore, by using linear basis functions and relatively small neighborhoods, this approach can approximate a large class of value functions, which may not be true for global approximations with any fixed set of basis functions. To store samples from all previous iterations and efficiently search for samples within a given neighborhood, we use a kd-tree data structure [7].

3 Model Validation

We first validate both seven-stage ADP implementations with no uncertainty in abatement cost against the results of the original DICE model, solved as a deterministic NLP. For ease of comparison, we modify the deterministic version of DICE to choose a single emissions control rate for each 50-year decision stage, rather than a distinct control rate for each decade. For the ADP version, we eliminate the uncertainty in abatement cost growth rates in order to approximate the deterministic solution. Figure 4 shows the resulting optimal decisions, as the fractional reduction below the reference emissions. The resulting optimal control from both ADP versions consistently converge to within a few percent of the NLP optimal controls. To test for convergence, we use the change in the moving average of the total Bellman value in the initial stage:

$$\bar{\epsilon} = |\mu_V^i - \mu_V^{i-1}|. \quad (6)$$

Because convergence is stochastic, we use a running average of 1000 sample values. The evolution of the convergence criterion over iterations is shown in Figure 5 for representative solutions of the global regression and the moving least squares algorithms. In general, the MLS algorithm converges in fewer iterations, typically 2000 iterations beyond the bootstrap for MLS as compared with over 10000 iterations for global regression to achieve a convergence criterion below $1e-7$. The tradeoff in computation speed is less clear, because each iteration of the MLS algorithm is more expensive, including augmenting the kd-tree, numerous searches of the kd-tree for nearest neighbors, and numerous regressions. The total computation time to achieve convergence to less than $1e-7$ is roughly equivalent between the two approaches, and the numerical results are comparable. For the remainder of this paper, we present only the results from the global regression version.

Next, we validate both ADP implementations for the stochastic case against the results of a stochastic dynamic program (SDP) version of DICE solved using a traditional backward induction algorithm. Given the need to iterate over discrete actions $\mu(t)$, cost growth shocks $\theta(t)$, and the DICE model's seven-dimensional state variable, we choose a two-stage model as the maximum-dimension model that can be solved using traditional numerical techniques within a reasonable timeframe and at an appropriate level of resolution. Both

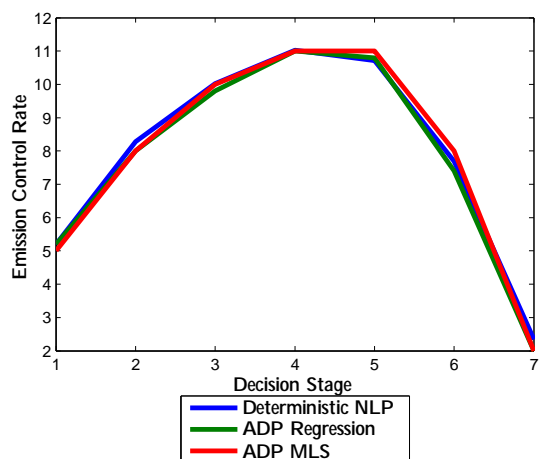


Fig. 4 Optimal control rate (%) in all decision stages when abatement cost is deterministic.




Fig. 5 Convergence criterion by iteration for global regression and moving least squares algorithms.

ADP methods converge to an optimal first-stage decision that is very close to that from backward induction, and the second stage optimal decisions are all in the same range, as shown in Table 1.

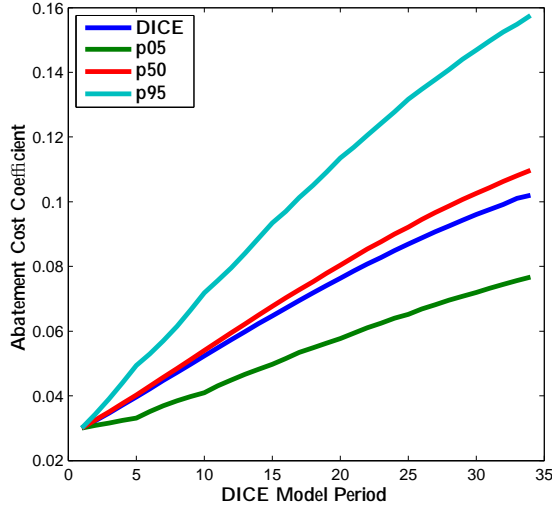
4 Results of Numerical Experiments

4.1 Results without Decision-Dependence

We now introduce uncertainty in the growth of the abatement cost coefficient as described above. We assume a reference value for the standard deviation in

Table 1 Optimal Controls for Stochastic Two-Stage Problem

Implementation	Stage 1	Stage 2 p5	p50	p95
Backward Induction	0.0705	0.0581	0.105	0.187
ADP Regression	0.0766	0.0412	0.115	0.219
ADP MLS	0.0757	0.0511	0.103	0.197

**Fig. 6** Uncertainty in abatement cost coefficient based on 20000 iterations and normally distributed i.i.d. multiplicative shock to growth rate with mean 1 and standard deviation 0.4.

cost shocks of 40% based on [32]. Figure 6 shows the resulting ranges in the abatement cost coefficient over time that results from 20000 samples in each stage of the normally distributed multiplicative shock applied to the growth rate. Using the ADP algorithm with global regression, we solve a seven-stage stochastic dynamic program with uncertain growth in abatement cost in each stage. As an example of the surface approximation of the Bellman value functions the stage 4 value function surface is shown in Figure 7. The surface for the other stages has the same general shape. The resulting optimal controls over the decision stages is shown in Figure 8. While the initial stage has a single optimal control rate, the optimal controls in all remaining stages are probabilistic and state-dependent. We therefore present the resulting optimal controls after convergence to the optimal policy using box plots. As a basis of comparison, we also plot the optimal controls from the deterministic NLP version of DICE. Note that the optimal first stage decision is a slightly higher control rate than the deterministic model. The same effect is seen in Figure 9, where an increase in the standard deviation of the cost shock while preserving the mean leads to slightly higher optimal first-stage controls (i.e., lower emissions). However, this effect of exogenous uncertainty on optimal control is quite small in magnitude.




Fig. 7 Representative surface of Bellman Value function by state variables, Capital Stock (K) and Temperature (TE).

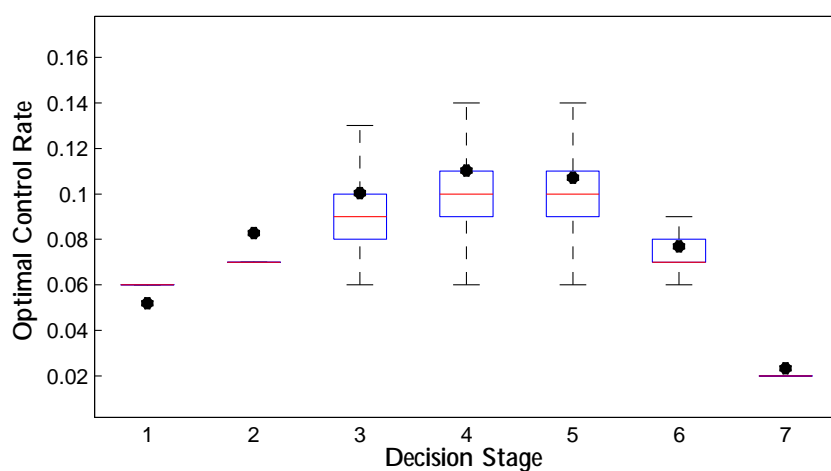


Fig. 8 Optimal control rate in 7-stage model with uncertain abatement costs. Box plots indicate the optimal control rates in stages 2-7, which are state-dependent. Black markers indicate the optimal control rates from a deterministic NLP version of DICE.

4.2 Results with Decision-Dependent Uncertainty

We now turn to the endogenous decision-dependent version of the model. As described in Section 2.2, the mean of the distribution of abatement costs in each stage is now a function of previous abatement decisions. This stylized model is intended to capture the generic effects of assuming that emissions limits induce technological change, which in turn lowers the costs of future abatement. Although not resolved in this model, this phenomenon could rep-

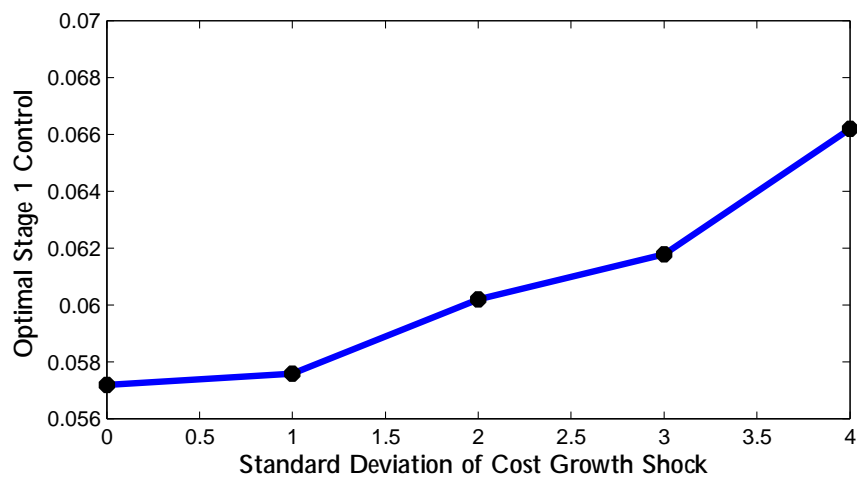


Fig. 9 Optimal emissions control rate in first decision stage as a function of the standard deviation of the cost growth rate shock.

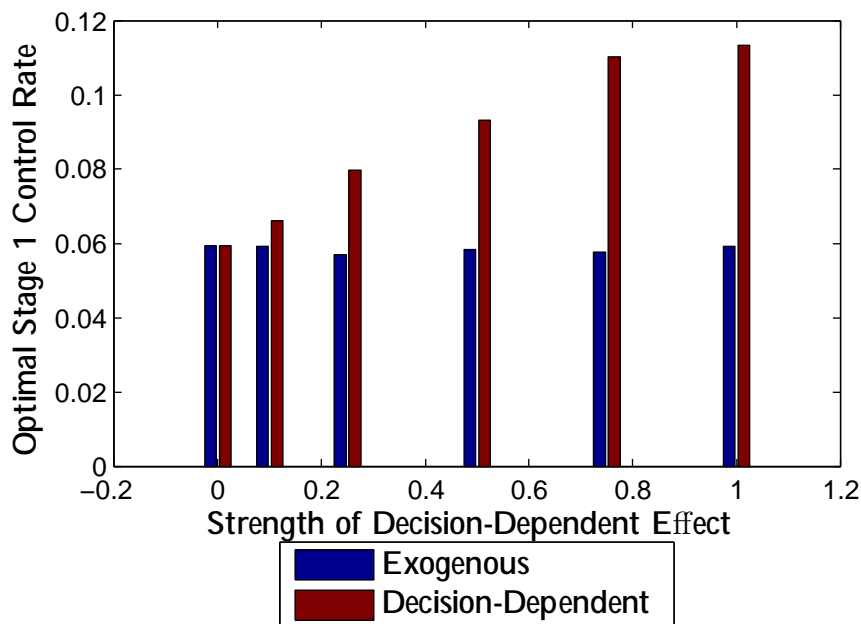


Fig. 10 Optimal Stage 1 control rate for different strengths of the decision-dependency.

resent either learning-by-doing in technological change [1,36,35] or induced R&D spending leading to technological change [23,11].

The stylized model used here is not the appropriate tool for determining the empirical strength of the decision-dependent effect, so we present a sensitivity analysis over a range of values for the strength of this effect (Figure




Fig. 11 Optimal Stage 1 control rate for different numbers of stages. Shown for case with exogenous cost uncertainty and decision-dependent cost uncertainty

10). Note that a value of zero for the parameter α is equivalent to the original reference model presented above, and a value of 1 indicates that the distribution of cost shocks almost entirely depends on the previous decision. To compare the effect of adding the decision-dependency of the cost uncertainty, we need an appropriate comparison, because the decision-dependent effect leads to different distributions of cost-uncertainty with lower means. For each value of the α parameter presented, we compare against a case with exogenous cost uncertainty, in which the abatement cost is drawn (i.i.d.) from the same distribution of costs that resulted under the decision-dependent version but independent of the control rate chosen in the previous stage. For even a relatively weak decision-dependent effect (e.g., $\alpha = 0.1$) a higher control rate is optimal in stage 1. In terms of the Bellman equation, there is additional marginal value in the next stage ($t + 1$) to each unit of abatement in period t . In contrast, the exogenous case with identical abatement cost distributions leads to the same optimal control rate in the first stage, since there is no added value in future stages to controlling more now. This source of additional value from decision-dependency, in terms of inducing technological change, is often a primary motivation for near-term emissions reductions. However, it is typically not represented within integrated assessment models except for specialized versions focused on developing representations for technical change.

A key question in the literature on sequential decision under uncertainty is whether the often-used simplification of two-stage models to approximate multi-stage finite-horizon problems is an appropriate and robust simplification. As discussed in the introduction, restricting analysis to two-stage models has been a common strategy for addressing the computational burden of traditional stochastic programming and stochastic dynamic programming methods. The ADP approach presented here allows many more stages to be represented and still converges to a solution within seconds to minutes. In Figure 11 We compare the optimal first stage control rates that result for finite-horizon problems with different numbers of total stages. We compare this for both the model with exogenous uncertainty in abatement costs and the model with moderate strength ($\alpha = 0.5$) decision-dependent cost uncertainty. When the uncertainty in abatement cost is exogenous to the control rates chosen, the effect of more decision stages is to reduce the optimal control rate in the initial period. In contrast, when cost uncertainty is influenced by the control rate chosen, the effect of more future decision stages is to increase the optimal initial control rate. Because the representation here assumes that the reductions in the mean of the cost distribution are permanent and can be compounded, more future stages increase the marginal value of controlling emissions in the current stage. This result suggests that for exogenous uncertainties, the traditional two-stage approximation may be relatively robust. In contrast, if the focus of analysis is on R&D or other path-dependent decisions, a two-stage model may not capture the full value of near-term investments in a multi-stage context.

5 Discussion

Decision problems about how to respond to climate change will occur gradually over very long time-scales, under a great deal of uncertainty, and with learning along the way about these uncertainties. Further, sociotechnical systems, such as the global economic-energy-emissions nexus, often exhibit feedbacks between decisions and uncertainties over time, in contrast to assumptions of exogeneity conventionally used to keep computations tractable. Use of optimization and simulation models of such complex systems typically neglect uncertainty or simplify the model to two decision stages, where uncertainty is resolved in the second period. To explore more realistic decision problems, we require algorithms that can tractably solve multi-period stochastic optimization with multi-dimensional state variables and feedbacks between decisions and uncertainties.

We have demonstrated two variants on such an algorithm here, and applied it to the problem of optimal emissions abatement under abatement cost uncertainty to respond to climate change. These algorithms converge quickly without having sacrificed resolution in state space, action space, or number of decision stages. Using these algorithms, we have demonstrated that 1) path-dependency can increase the optimal level of first stage controls, and 2) this effect is increasing in the number of decision stages modeled. This suggests

that treating decision-dependent phenomenon as exogenous and two-stage approximations of multi-stage problems where there is decision-dependency could result in bias towards lower first stage controls. Such a bias would result from neglecting a component of the true value in future states of near-term investments and controls. The example here is merely suggestive, and exploration of other uncertainties as well as theoretical developments are needed to more fully explore this question.

We have also compared two alternative methods for value function approximation, one parametric and one non-parametric. The non-parametric approach using moving least squares converges in fewer iterations, but each iteration is more computationally expensive. This particular application happens to have smooth value function surfaces that are well approximated by global second-order polynomials. However, other questions may involve highly nonlinear or discontinuous surfaces, such as considering threshold effects or tipping points in the climate system. We expect that the non-parametric approach will have significant advantages for such applications, and constitutes a basis of our future work with these algorithms.

Acknowledgements The authors would like to thank David Popp, Karen Fisher-Vanden, Richard Larson, John Parsons, and two anonymous reviewers for helpful comments. We gratefully acknowledge the support of the U.S. National Science Foundation grant number 0825915 and the U.S. Department of Energy grant number DE-SC0003906.

References

1. Arrow, K.: The economic implications of learning by doing. *Review of Economic Studies* **29**, 155–173 (1962)
2. Baker, E., Solak, S.: Optimal Climate Change Policy: R&D Investments and Abatement under Uncertainty (2011). Under Review
3. Bellman, R.: *Dynamic programming*. Dover Pubns (2003)
4. Bertsekas, D.: *Dynamic programming and optimal control*. Athena Scientific Belmont, MA (2007)
5. Bertsekas, D., Tsitsiklis, J.: *Neuro dynamic programming*. Athena Scientific Belmont, MA (1996)
6. Crost, B., Traeger, C.: Risk and aversion in the integrated assessment of climate change (2010). CUDARE Working Paper No. 1104
7. Fasshauer, G.: *Meshfree approximation methods with Matlab*. World Scientific Pub Co Inc (2007)
8. Gerst, M., Howarth, R., Borsuk, M.: Accounting for the risk of extreme outcomes in an integrated assessment of climate change. *Energy Policy* **38**(8), 4540–4548 (2010)
9. Goel, V., Grossmann, I.: A class of stochastic programs with decision dependent uncertainty. *Mathematical programming* **108**(2), 355–394 (2006)
10. Hammitt, J.K., Lempert, R.A., Schlesinger, M.E.: A sequential-decision strategy for abating climate change. *Nature* **357**, 315–318 (1992)
11. Jaffe, A., Newell, R., Stevens, R.: *Handbook of Environmental Economics, Vol 1.*, chap. Technological change and the environment, pp. 461–516. North-Holland/Elsevier (2003)
12. Kelly, D., Kolstad, C.: Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* **23**, 491–518 (1999)
13. Kolstad, C.: Learning and stock effects in environmental regulation: The case of greenhouse gas emissions. *Journal of Environmental Economics and Management* **31**, 1–18 (1996)

14. Leach, A.: The climate change learning curve. *Journal of Economic Dynamics and Control* **31**, 1728–1752 (2007)
15. Lemoine, D., Traeger, C.: Tipping points and ambiguity in the integrated assessment of climate change (2011). NBER Environmental and Energy Economics Summer Institute 2011 Paper
16. Manne, A., Richels, R.: The costs of stabilizing global co2 emissions: A probabilistic analysis based on expert judgment. *The Energy Journal* **15**(1), 31–56 (1994)
17. Martens, P., Rotmans, J.: Climate change: An integrated perspective. In: P. Martens, J. Rotmans, D. Jansen, K. Vrieze (eds.) *Climate Change: An Integrated Perspective, Advances in Global Change Research*, vol. 1, pp. 1–10. Springer Netherlands (2003)
18. McKay, M.D., Beckman, R.J., Conover, W.J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**(2), 239–245 (1979)
19. Nordhaus, W.: The challenge of global warming: Economic models and environmental policy. Available at: <http://nordhaus.econ.yale.edu/> (2007). NBER Working Paper 14832
20. Nordhaus, W., Boyer, J.: *Warming the world: economic modeling of global warming*. Massachusetts Institute of Technology Press, Cambridge, MA, USA (2000)
21. Nordhaus, W., Popp, D.: What is the value of scientific knowledge? an application to global warming using the price model. *The Energy Journal* **18**(1), 1–45 (1997)
22. Parpas, P., Webster, M.: A Stochastic Minimum Principle and a mesh-free method for Stochastic Optimal Control (2011). Submitted
23. Popp, D., Newell, R., Jaffe, A.: *Energy, the environment, and technological change*. howpublished (2009). NBER Working Paper 14832
24. Popp D., S.N.F.V.K., Webster, M.: Technology variation vs. r&d uncertainty: What matters most for energy patent success? howpublished (2011). NBER Working Paper No. 17792
25. Powell, W.: *Approximate Dynamic Programming: Solving the curses of dimensionality*, vol. 703. Wiley-Blackwell (2007)
26. Reilly, J., Edmonds, J., Gardner, R., Brenkert, A.: Monte carlo analysis of the iea/orau energy/carbon emissions model. *The Energy Journal* **8**(3), 1–29 (1987)
27. Scott, M., Sands, R., Edmonds, J., Liebetrau, A., Engel, D.: Uncertainty in integrated assessment models: Modeling with minicam 1.0. *Energy Policy* **27**(14), 597–603 (1999)
28. Stock, J.H., Watson, M.W.: Variable trends in economic time series. *The Journal of Economic Perspectives* **2**(3), 147–174 (1988)
29. Webster, M.: The curious role of learning: Should we wait for more data? *The Energy Journal* **23**, 97–119 (2002)
30. Webster, M.: Incorporating path-dependency into decision analytic methods: An application to global climate change policy. *Decision Analysis* **5**(2), 60–75 (2008)
31. Webster, M., Jakobovits, L., Norton, J.: Learning about climate change and implications for near-term policy. *Climatic Change* **89**(1-2), 67–85 (2008)
32. Webster, M., Paltsev, S., Parsons, J., Reilly, J., Jacoby, H.: Uncertainty in greenhouse emissions and costs of atmospheric stabilization. Tech. rep., MIT JPSPGC (2008). Report No. 165
33. Webster, M., Sokolov, A., Reilly, J., Forest, C., Paltsev, S., Schlosser, A., Wang, C., Kicklighter, D., Sarofim, M., Melillo, J., Prinn, R., Jacoby, H.: Analysis of climate policy targets under uncertainty. Tech. rep., MIT JPSPGC (2009). Report No. 180
34. Weyant, J., Davidson, O., Dowlabathi, H., Edmonds, J., Grubb, M., Parson, E., Richels, R., Rotmans, J., Shukla, P., Tol, R., Cline, W., Fankhauser, S.: Integrated assessment of climate change: an overview and comparison of approaches and results. In: *Economic and Social Dimensions of Climate Change, Equity and Social Considerations*. Contribution of Working Group III to the Second Assessment Report of the Intergovernmental Panel on Climate Change, chap. Integrated assessment of climate change: an overview and comparison of approaches and results, pp. 367–396. Cambridge University Press (1996)
35. Wright, T.: Factors affecting the cost of airplanes. *Journal of Aeronautical Sciences* **3**, 122–128 (1936)
36. Wright, T.: *Inside the Black Box: Technology and Economics*. Cambridge University Press (1982)
37. Yohe, G., Andronova, N., Schlesinger, M.: To hedge or not against an uncertain climate future? *Science* **306**, 416–417 (2004)

Appendix

In this Appendix, we present the DICE model equations.

$$\max_{I(t), \mu(t)} W = \sum_t U[c(t), L(t)]R(t). \quad (\text{A.1})$$

$$R(t) = \prod_{\nu=0}^t [1 + \rho(\nu)]^{-10} \quad (\text{A.2})$$

$$\rho(t) = \rho(0)e^{-g^{\rho}t}.$$

$$U[c(t), L(t)] = L(t)\log[c(t)]. \quad (\text{A.3})$$

$$g^{pop}(t) = g^{pop}(0)e^{-\delta^{pop}t} \quad (\text{A.4})$$

$$L(t) = L(0)e^{\int_0^t g^{pop}t}$$

$$Y(t) = \Omega(t)(1 - b_1(t)\mu(t)^{b_2})A(t)K(t)^{\gamma}L(t)^{1-\gamma} \quad (\text{A.5})$$

$$g^A(t) = g^A(0)e^{-\delta^A t} \quad (\text{A.6})$$

$$A(t) = A(0)e^{\int_0^t g^A t}$$

$$\Omega(t) = 1/[1 + D(t)] \quad (\text{A.7})$$

$$D(t) = \theta_1 T(t) + \theta_2 T(t)^2 \quad (\text{A.8})$$

$$g^b(t) = g^b(0)e^{-\delta^b t} \quad (\text{A.9})$$

$$b_1(t) = b_1(t-1)/(1 + g^b(t))$$

$$b_1(0) = b_1^*$$

$$E(t) = (1 - \mu(t))\sigma(t)A(t)K(t)^{\gamma}L(t)^{1-\gamma} \quad (\text{A.10})$$

$$g^{\sigma}(t) = g^{\sigma}(0)e^{-\delta_1^{\sigma}t - \delta_2^{\sigma}t^2} \quad (\text{A.11})$$

$$\sigma(t) = \sigma(t-1)/(1 + g^{\sigma}(t))$$

$$\sigma(0) = \sigma^*$$

$$Y(t) = C(t) + I(t) \quad (\text{A.12})$$

$$c(t) = C(t)/L(t) \quad (\text{A.13})$$

$$K(t) = (1 - \delta_K)K(t-1) + 10 \times I(t-1) \quad (\text{A.14})$$

$$K(0) = K^*$$

$$LU(t) = LU(0)(1 - \delta_1 t) \quad (\text{A.15})$$

$$ET(t) = E(t) + LU(t)$$

$$M_{AT}(t) = 10 \times ET(t-1) + \phi_{11}M_{AT}(t-1) - \phi_{12}M_{AT}(t-1) + \phi_{21}M_{UP}(t-1) \quad (\text{A.16})$$

$$M_{AT}(0) = M_{AT}^*$$

$$M_{UP}(t) = \phi_{22}M_{UP}(t-1) + \phi_{12}M_{AT}(t-1) + \phi_{32}M_{LO}(t-1) \quad (\text{A.17})$$

$$M_{UP}(0) = M_{UP}^*$$

$$M_{LO}(t) = \phi_{33}M_{LO}(t-1) + \phi_{23}M_{UP}(t-1) \quad (\text{A.18})$$

$$M_{LO}(0) = M_{LO}^*$$

$$F(t) = \eta \log[M_{AT}(t)/M_{AT}^{PI}] / \log(2) + O(t) \quad (\text{A.19})$$

$$\begin{aligned} O(t) &= -0.1965 + 0.13465t & t \leq 11 \\ &= 1.15 & t > 11 \end{aligned}$$

$$T(t) = T(t-1) + \sigma_1 F(t) - \lambda T(t-1) - \sigma_2 [T(t-1) - T_{LO}(t-1)] \quad (\text{A.20})$$

$$T(0) = T^*$$

$$T_{LO}(t) = T_{LO}(t-1) + \sigma_1 F(t) - \lambda T(t-1) - \sigma_2 [T(t-1) - T_{LO}(t-1)] \quad (\text{A.21})$$

$$T_{LO}(0) = T_{LO}^*$$

Table 2 Nonlinear Model Results

Parameter	Value	Description
$\rho(0)$	0.03	Initial rate of social time preference
g^p	0.25719	Decline rate of social time preference
$g^{pop}(0)$	0	a
δ^{pop}	0	a
$L(0)$	0	a
b_2	0	a
γ	0	a
a	0	a
a	0	a
a	0	a
a	0	a
a	0	a
a	0	a
θ_1	-0.0045	Damage Coefficient Linear Term
θ_2	0.0035	Damage Coefficient Quadratic Term
$g^b(0)$	-0.08	Decline in Abatement Cost Function
δ^{ab}	0.5	Change in Decline of Cost Function
b_1^*	0.03	Intercept Control Costs Function