# The Climate Change Learning Curve

by Andrew LEACH

Cahier de recherche nº IEA-04-03 April 2004

ISSN: 0825-8643

Copyright © 2004 HEC Montréal.

Direction de la recherche, HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine, Montréal (Québec) Canada H3T 2A7.

Tous droits réservés pour tous pays. Toute traduction ou toute reproduction sous quelque forme que ce soit est interdite. Les textes publiés dans la série des Cahiers de recherche HEC n'engagent que la responsabilité de leurs auteurs. La publication de ce Cahier de recherche a été rendue possible grâce à des subventions d'aide à la publication et à la diffusion de la recherche provenant des fonds de l'École des HEC.

# The Climate Change Learning Curve

Andrew J. Leach<sup>\*</sup> Institut d'économie appliquée

HEC Montréal

This Version: April 22, 2004

#### Abstract

The key element in the tension between those who believe climate change is an issue and those who do not is essentially the question of whether we are merely in a long period of shock-induced above average temperatures or if we have led to this increase in temperatures by anthropogenic carbon emissions. The model proposed in this paper allows for a model in which we weigh observations on temperature against the potential that these are generated by a combination of uncertain parameters; namely the coefficient of autoregression and the sensitivity of temperature change to atmospheric carbon levels. This paper shows that, contrary to predictions in the literature that we can resolve uncertainty very quickly, the time to learn may be on the order of thousands of years when uncertainty surrounds two parameters in the law of motion for temperature. When the learning model is embedded in an optimal policy growth model, policy decisions are found to be affected by the prior mean but not the variance. A new solution algorithm which relies on randomization and least squares approximation is applied to solve the value function in the model.

JEL classification: Q25; Q28; D83; D81; C61; C63; E1; E61; H4

*Keywords*: Climate Change; Bayesian Learning; Environmental Regulation; Growth; Pollution; Dynamic Programming; Precautionary Principle.

<sup>\*</sup>I would like to thank Chris Ferrall, Kim Huynh, Geoffrey Dunbar, Bernard Sinclair-Desgagné, Ross McKitrick and Michelle Reinsborough who provided helpful discussion. Thanks are also due to Jeremy Lise and Jean-François Houde who provided parallel Ox code. This research was completed in part using the resources of the High Performance Computing Virtual Laboratory at Queen's University. Thanks to conference participants at the Canadian Resource and Environmental Economics Study Group, 2003, and seminar participants at Université Laval and the 2004 CIREQ Resource and Environmental Economics Workshop for helpful comments. The usual disclaimer applies. Address communications to the author at andrew.leach@hec.ca or Institut d'économie appliquée, HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine, Montréal (Québec) H3T 2A7.

# 1 Introduction

The tension between those who believe climate change is an issue and those who do not is summarized by the question of whether we are merely in a long period of naturally occurring, above average temperatures or if we have led to this increase in temperatures through anthropogenic carbon emissions. Of course, the case of either can be strengthened by arguing in favor of long-lived shocks to temperature. Kondratyev (1988) states that taking as given that temperature shocks are long-lived, we could argue either that we are in a long period of warming, and the effect of carbon is minimal, or, alternatively, that we are in a long period of cooling, but these natural changes have been reversed by anthropogenic carbon emissions and when the trend changes we will be in a very serious situation. Either way, it is essential to note that, rather than seeing temperature as having a fixed autoregressive relationship and learning about the effect of carbon, agents may well weigh observations on temperature against the potential that these are generated by a combination of uncertain parameters; namely the coefficient of autoregression and the sensitivity of temperature change to atmospheric carbon levels.

Consideration of this type of uncertainty is of paramount importance in addressing the problem of climate change and global mitigation policies. A literature has emerged which confronts this challenge to policy makers through the examination of optimal policy under uncertainty. To this point, the literature has not addressed the question of how to confront the problem that uncertainty over climate change occurs on many parameters simultaneously. This paper seeks to address this dimension by exploring the problem of two-dimensional uncertainty. Scientific observation has informed us that temperature evolves in a stochastic manner, with long periods of below or aboveaverage temperatures. At the same time, there is scientific uncertainty about the parameters which determine the influence of carbon emissions on temperature changes. Uncertainty affecting these two possibilities simultaneously is important since the effects of the two cannot be easily distinguished. The model proposed in this paper captures the inability to efficiently distinguish from the data cyclical warming versus carbon-induced climate changes. This paper deals with two key questions: first, how can we expect uncertainty over the parameters governing climate change to evolve over time, and second, how does uncertainty and learning affect the evolution of optimal policy.

This paper first seeks to explore, in a reduced form environment, the dynamics of learning in a climate change model. A previous paper by Kelly and Kolstad (1999a) evaluates the expected time to resolve uncertainty about the value of climate sensitivity to  $CO_2$  levels when only the parameter

of interest is not known with certainty. When the autoregressive law of motion for temperature is also uncertain, and two parameters must be estimated simultaneously, uncertainty is expected to be substantially more persistent. A Monte Carlo experiment is presented to demonstrate expected learning times, and provides evidence of a strong effect on learning times.

The second contribution of this paper is to examine the effect of parametric uncertainty in an optimal policy model with learning. The learning model used in the first section is embedded in an adapted Nordhaus (1994) model of climate and economy. The model is calibrated, and simulations are used to provide predictions on two effects. First, we want to know how optimal carbon emissions control changes under different levels of uncertainty. Consider a regulator who is faced with the current policy environment: a recent trend of higher temperatures combined with levels of atmospheric greenhouse gases that are higher than ever. The regulator can essentially conclude that the temperature changes are the result of natural, long-lived shocks and that the effect of carbon is minimal and thus hold off on regulation, or they can decide that the change is carbon-induced, and immediately institute regulations to curb emissions. The effect we wish to capture is, in a model with endogenous emissions, how persistent is parametric uncertainty and how does it impact policy choices.

The results of the paper show that learning times are very significantly affected by the number of parameters over which there is uncertainty. If the average growth rate of emissions is assumed to be 2.5% per year, Monte Carlo experiments find that while it would take 72.5 years to reach relative certainty about a single uncertain parameter with continuous domain, adding uncertainty about a second parameter extends this time to 220 years. Learning about two parameters is introduced to an optimal control model, and the effect on optimal control levels is marginal, and depends more on the estimated parameter value than on the uncertainty surrounding it. The simulations used allow for the identification of the magnitude of precaution employed in an optimal regulation framework. When emissions are endogenous in a model of the global economy, learning times are further lengthened by emissions reduction activity in each time period, which reduces the variance of emissions.

A second important result reveals the fact that, due to the inability to separately identify the effects of natural and anthropogenic forcing, there is a great deal more importance on the accuracy of the initial prior. Essentially, the paper shows a result of self-fulfilling expectations where if we believe that natural climate cycles are responsible for temperature fluctuations, and we take this prior to the data, our prior error will continue to influence the learned estimate until we stop

emitting greenhouse gases altogether.

This paper proceeds as follows. Section 2 presents a review of the relevant literature on learning and climate change economics and policy. Section 3 discusses the Bayesian learning approach used in this paper and presents the Monte Carlo learning experiment. Section 4 presents the economic learning model. Section 5 presents the solution algorithm, and Section 6 presents simulations and optimal policy results. Section 7 concludes.

## 2 Relevant Literature

In order to place this work in the context of the literature on learning and climate change, a brief summary of some of the relevant papers is provided below. The benchmark works on climate change economics are Manne and Richels (1992), Manne, Mendelsohn and Richels (1995), Nordhaus (1994), and Nordhaus and Boyer (2000). These works all contain extensive reference to uncertainty. A related literature on learning and irreversibility in climate change includes papers by Kolstad (1994, 1996, 1997), Kelly and Kolstad (1999a), Kelly, Kolstad, and Mitchell (1999), and Ulph and Ulph (1997). The application of Bayesian learning to economic models is pioneered in an article by Cyert and DeGroot (1974).

Nordhaus (1994) and, more recently, Nordhaus and Boyer (2000) present a dynamic, regionally disaggregated model of global climate and economy. Their model is deterministic and seeks to derive the optimal abatement policy as a solution to a finite horizon Ramsey (1928) problem. While neither book specifically models uncertainty and information, the topic is never far from the surface. Specifically, Nordhaus and Boyer discuss the range of estimates currently in the literature for many of the climate parameters which are fixed in their simulations. They discuss sensitivity to the rate of radiative forcing from an increase in atmospheric  $CO_2$  concentrations, the damages due to temperature increase and to a lesser degree, the mechanism of temperature change itself.

Manne and Richels (1992) and, more recently, Manne, Mendelsohn and Richels (1995) present a model for evaluating regional and global effects of greenhouse gas reduction policies. Their 1992 paper begins with a discussion of the uncertainties over the costs of reducing emissions. They also discuss directly the uncertainty over the level to which increasing concentrations of greenhouse gases in the atmosphere will affect climate in the future, and discuss the reduction of these damages to a dollar figure. Manne and Richels (1992) presents a treatment of learning in an attempt to answer the question of whether it makes sense to learn then act, or vice versa. Learning in their model is exogenous and occurs completely after three time periods. Their economy begins in 1990, where the regulator chooses to act then learn or vice-versa. In the act-then-learn scenario, once learning is complete, the economy follows an optimal path. Their paper also addresses the question of the value of perfect information, since they can compare the outcome of the economy in a learning framework to a benchmark, perfect information case.

Pizer (1996) presents an analysis in which the effects of parametric uncertainty are examined within an optimal policy model. In this paper, data are used to estimate distributions of parameters in a stochastic optimal growth model. The role of this parametric uncertainty is examined to demonstrate the effect of not only the uncertain economic state, but also uncertainty over elements in the laws of motion in the economy. It is this type of uncertainty that permeates climate change. Global temperature and atmospheric carbon dioxide are both stocks with very slow rates of change. For this reason, there is little effect from shocks on an annual basis. Conversely, the effect of a small change in the parameters governing the half-life of the stock can be very significant.

In Kolstad(1996), the tension between acting now or postponing until we have more information is explored in an environment with some irreversibility. The irreversibility lies in the fact that emissions are assumed to be strictly positive, such that carbon concentrations in the atmosphere can only be reduced through natural decay. The duration of the stock effect is thus largely a function of the long half-life of carbon in the atmosphere. Climate change, if it exists, is not caused by the flow of emissions, but rather the existing stock of  $CO_2$  and other greenhouse gases in the atmosphere. Emissions today contribute to that stock, and decay at a very slow rate (half life of 80-120 years), so they are irreversible from the perspective of agents alive today. The level of capital devoted to abatement may also be irreversible. The Kolstad (1996) model examines an environment where a regulator learns about the level of damages from temperature changes, and can adapt policies over time. In this environment, an easily reversible policy such as a tax is preferred to a level of abatement capital which is non-reversible in the short run.

Ulph and Ulph (1997) examine a similar question to Kolstad (1996): how the potential to acquire better information about damages from climate change in the future affects today's policy decisions. Specifically, they look at the application of Epstein's (1980) results on optimal decision making in the face of irreversibility to the problem of climate change. Intuitively, if an action is irreversible, the possibility of acquiring better information in the future should lead to a lower level of commitment today. Their paper finds several important results from the addition of learning and uncertainty to a Maddison (1995) model of climate change. They show that the irreversibility

effect does not hold for all cases of climate change, and that there is little empirical support for such an effect. Furthermore, they find that as the discount factor is decreased and/or uncertainty is increased, the effect of learning on the optimal decision time frame is increased.

Kelly and Kolstad (1999a) present the paper which is most closely related to the exercise undertaken here. They construct a model in which a social planner uses information from temperature realizations to update prior beliefs about the sensitivity of climate (temperature) to atmospheric greenhouse gas levels. In the model, the planner chooses the optimal level of savings and emissions control conditional on their knowledge of the mechanism of climate change at each point in time, updates those beliefs, and thus adjusts their actions, conditional on observations from the climate record. Learning is Bayesian, thus the planner is using information in an optimal manner. A key result in their model is that the expected learning time is 90-160 years, which is the time after which parametric uncertainty is essentially removed from the planner's problem. Kelly and Kolstad (1999a) apply the computational techniques developed in Kelly and Kolstad (1999b) to solve the stochastic climate and economy model, and to simulate the results.

# 3 Learning about Multiple Parameters

The goal of this section of the paper is to introduce an experiment which captures the essence of the problem with resolving uncertainty about the effect of  $CO_2$  on climate. This experiment is meant to capture the learning dynamics discussed in the Introduction, where conflicting hypotheses may be supported with equal likelihood by the data.

#### 3.1 Learning Model

Global temperature is characterized by long periods of warming and cooling, which is assumed to be driven by a first order vector autoregressive process. Let G represent global surface temperature deviations in degrees Celcius, measured relative to the pre-industrial global average, while Orepresents deviations in temperature in the world's upper oceans, globally and seasonally averaged. For example, for the year 1985, Nordhaus and Boyer (2000) use starting values for these variables of .43°C and .06°C respectively. Surface temperature is modelled as a function of previous period levels of surface and ocean temperature, as well as the relative concentration of atmospheric carbon dioxide relative to pre-industrial levels  $\left(\frac{m_t}{m_b}\right)$ .<sup>1</sup>

$$G_{t+1} = \lambda_1 G_t + \omega O_t + \eta \frac{\log \frac{m_t}{m_b}}{\log 2} + u_t \tag{1}$$

such that as atmospheric carbon levels  $m_t$  limit to preindustrial levels, temperatures will tend to zero. The parameter  $\lambda_1 \in (0,1)$  is the coefficient of autoregression<sup>2</sup>.  $\eta$  governs the climate sensitivity to carbon concentration, and captures the long-run effect on surface temperature in  ${}^{o}C$  of a long-run doubling of  $CO_2{}^3$ .  $\omega$  measures the effect of ocean temperature deviations on surface temperature deviations, which captures the buffering effect of ocean temperatures.  $u_t$  is a normally and identically distributed shock with mean 0 and variance  $\sigma_u^2$ . The pre-industrial atmospheric carbon mass is captured by the parameter  $m_b$ , such that increases in atmospheric  $CO_2$  are expressed in relative terms.

The evolution of temperature, or climate change, occurs through a slow warming of the world's oceans and atmosphere, which is prevented in the short run by thermal inertia. The process is buffered by slow changing ocean temperatures which are modelled as an autoregressive process with parameter  $\lambda_2 \in (0, 1)$ :

$$O_{t+1} = \lambda_2 O_t + (1 - \lambda_2) G_t \tag{2}$$

In order to simplify notation for the learning model, define  $H_t \equiv G_{t+1} - \omega O_t$  and  $X_t \equiv \{G_t, \frac{\log \frac{m_t}{m_b}}{\log 2}\}$ . Also, let the learning parameter vector be  $\mu \equiv (\hat{\lambda}_1 \ \hat{\eta})'$ .

The learning model employed here is Bayesian, and takes as given  $\sigma_u^2$ , the variance of the shock in the law of motion for surface temperature. Uncertainty exists over the values of  $\lambda_1$  and  $\eta$ , while parameters  $\lambda_2$  and  $\omega$  are assumed known with certainty<sup>4</sup>. Let the prior distribution be a multi-variate normal with  $2 \times 1$  vector of mean estimates  $\mu$  and  $2 \times 2$  variance-covariance matrix **V**. Updating according to observations on **H** and **X** in each time period leads to a posterior distribution which is normal with mean  $\mu^*$  and variance-covariance matrix **V**<sup>\*</sup>. The updating rules

 $<sup>{}^{1}\</sup>text{CO}_{2}$  may be measure in concentration (ppmv) or mass (GtC), as long as  $m_{t}$  and  $m_{b}$  are measured in the same units. The conversion factor is 1ppmV=2.13 GtC.

<sup>&</sup>lt;sup>2</sup>There exists some dispute over the existence of a unit root in temperature. Using an annual measurement interval, data from the CDIAC was tested and an Augmented Dickey Fuller test rejects the null of a unit root at the 5% level

<sup>&</sup>lt;sup>3</sup>This relationship is operationalized as in Nordhaus and Boyer (2000) in order to provide an interpretable unit for the coefficient  $\eta$ .

<sup>&</sup>lt;sup>4</sup>The effects shown below will be magnified as uncertainty is added over additional parameters.

are given as follows:

$$\mu^* = \left[ \mathbf{V}^{-1} + \left[ \sigma_u^2 (X^T X) \right]^{-1} \right]^{-1} \left[ \mathbf{V}^{-1} \mu + \left[ \sigma_u^2 (X^T X)^{-1} \right]^{-1} (X^T X)^{-1} X^T H \right]$$
(3)  
=  $F \mu + (I - F) (X^T X)^{-1} X^T H$ (4)

where

$$F = \left[ \mathbf{V}^{-1} + \left[ \sigma_u^2 (X^T X) \right]^{-1} \right]^{-1} \mathbf{V}^{-1}$$
(5)

$$\mathbf{V}^{*} = \left[\mathbf{V}^{-1} + \left[\sigma_{u}^{2}(X^{T}X)^{-1}\right]^{-1}\right]^{-1}$$
(6)

where it is evident in (4) that the next estimate of the mean vector is a weighted average of the mean of the prior distribution and the standard OLS regression coefficient vector for the observations used for updating. The diagonal elements of the variance/covariance matrix will be monotonically decreasing with each observation.

#### **3.2** Monte Carlo Experiment

In Kelly and Kolstad (1999a), an analytic result is presented to describe the expected number of periods it takes for agents to reach a reasonable estimate of the true parameter governing a climate change process. In this section, a Monte Carlo experiment is undertaken to derive the expected learning curve. The experiment looks at the relationship between emissions growth rates and learning. The reason for this comparison is that the regulator is assumed to be able to control emissions to extract information about the mechanism of climate change.

**Experiment 1** : Learning time versus number of uncertain parameters and emissions growth  $rates(\gamma_E)$ .

Assume the following:

1. Atmospheric carbon concentration is determined by exogenous carbon emissions and stock decay ( $\delta_m = 0.99167$ ). Emissions and atmospheric carbon mass, each measured in gigatons of carbon (GtC), grow according to the following laws of motion:

$$E_t = E_{t-1} * (1 + (\gamma_E * .995^t))$$
  

$$m_t = E_t + (1 - \delta_m)(m_t - m_b) + m_b$$

2. Temperature evolves according to the two equation system given in (1-2).

3. The prior distribution on the parameters is defined by mean vector  $\mu \equiv (\hat{\lambda}_1 \ \hat{\eta})'$  and variance-covariance matrix **V**. Initial prior distribution has  $\mu \equiv \{.9122, 2\}$  such that the mean of the initial prior is the true value for the autoregressive parameter. The initial variances are chosen such that the confidence intervals are  $\mu \pm \{.1, 2\}$ .



Figure 1: Learning Time vs. Growth Rate of Emissions

- 4. Based on observing data  $\{H_t, X_t\}_{t=1..T}$ , and known parameter  $\sigma_u^2$ , beliefs are updated according to (3-6).
- 5. Learning has occurred when a value of  $\eta$  which is  $1^{\circ}$  away from the true parameter can be rejected at a 99% confidence level.

The expected learning times based on 3000 replications as a function of initial emissions growth rate are shown on the graph in Figure 1. The learning times when two parameters are uncertain exceed the single parameter learning times by an significant multiplier for all feasible emissions growth rates.

It is plain to see that the magnitude of learning times initially cited by Kelly and Kolstad (1999a) of less than 25 years increase greatly with inclusion of learning over two parameters. In this case, to generate learning times on the order of 25 years, growth rates of emissions need to be

over 50% per year over the time period<sup>5</sup>.



Figure 2: Prior Mean and Confidence Intervals over Time

It is important to note that there is a great deal of acceleration in the learning time generated by increases in emissions growth rates. This occurs because the induced variance in atmospheric carbon levels allows for a quicker identification of the parameter value in question. When emissions are increased, this directly affects atmospheric carbon levels. Decreasing emissions will also affect carbon levels, but not to the same degree, because the stock of carbon is long lived. With a halflife on the order of 70 years, reducing emissions (even ceasing emissions altogether) in any period will not generate significant variation in observed carbon levels. This result is very important if optimal policy models with learning are to be taken seriously. The cost of uncertainty, particularly in climate change regulation, may be very high. As such, if it is possible to raise emissions levels in the short term to reduce that uncertainty very quickly, it may appear to pay off to do so in

 $<sup>^{5}</sup>$ Over the period of 1960-1998, emissions growth rates averaged 2.5%

certain frameworks (ie the cost of increasing emissions may be less than the benefit of reducing uncertainty). When two or more sources of uncertainty are considered, it is less likely that the regulator will be able to increase welfare levels by requiring a short-lived spike in emissions because the time to learn will still be significant.

To appreciate the way learning occurs over time, consider one example from the above simulation, where emissions grow initially at 1.5% per year. In Figure 2, the curves show the prior mean and 99% confidence interval for the radiative forcing parameter  $\eta$ , along with a horizontal line at the true parameter value ( $\eta = 4$ ).

### 4 Integrated Assessment Model

This section embeds the learning characterization outlined above in a modified Nordhaus (1994) integrated assessment model (IAM) of the global economy. As compared to the learning experiment in the previous section, emissions are an endogenous result of economic activity and regulatory control. A social planner can chooses savings and emissions control rates in each period in order to maximize the sum of per capita utility. A benchmark model where the regulator faces only uncertainty over the stochastic realization of temperature is compared with a model where parametric uncertainty exists in this law of motion and diminishes over time with learning. This allows the examination of learning times with endogenous emissions generation and control, and the level to which this added uncertainty affects optimal policy.

#### 4.1 The Economic Environment

An initial population of  $L_0$  agents identical in all aspects grows through time at a convergent rate, such that population growth eventually limits to zero.

$$L_t = L_0 \, \exp\left(\frac{\gamma_l}{\delta_l} (1 - e^{-\delta_l t})\right) \tag{7}$$

Technological change in the model is exogenous, and also converges to a constant Hicks-neutral technology parameter A.

$$A_t = A_0 \exp\left(\frac{\gamma_a}{\delta_a}(1 - e^{-\delta_a t})\right).$$
(8)

Production of the aggregate consumption and investment good occurs according to a Cobb-

Douglas production function given by:

$$Y_t = \frac{1 - b_1 \tau_t^{b_2}}{1 + \theta_1 G^{\theta_2}} A_t K_t^{\alpha} L_t^{1-\alpha} = C_t + I_t$$
(9)

The equation includes the cost of CO<sub>2</sub> emissions control rate  $\tau \in [0, 1)$  and the cost of increased surface temperature G on the productivity of factors K and L. The parameters  $b_1$  and  $b_2$  are the linear and exponential components of regulatory costs, while  $\theta_1$  and  $\theta_2$  have the equivalent interpretation for damages from temperature deviations. This characterization of the relationship between production, regulation, and damages originates in Nordhaus (1994). In this characterization, abatement is completely reversible from period to period.  $\lambda$  is the standard Cobb-Douglas parameter. Production is assumed usable for both investment and consumption.

Capital stock (K) evolves endogenously with investment  $(I_t)$  and depreciates according to parameter  $(\delta_k \in [0, 1])$ :

$$K_{t+1} = (1 - \delta_k)K_t + I_t \tag{10}$$

Agents in the economy have a constant relative risk aversion utility function with coefficient of relative risk aversion  $\sigma^6$  given by:

$$U(C_t) = \frac{C_t^{1-\sigma}}{1-\sigma} \tag{11}$$

A social planner seeks to maximize discounted social welfare, where welfare in each period is defined as the sum of utility per capita.

#### 4.2 The Physical Environment

The role of the climate change model is to provide a law of motion for the climatic state as a function of the endogenous emissions of greenhouse gases. The schematic diagram in Figure 3 shows an overview of the system.

Emissions  $E_t$  into the atmosphere are governed by the exogenous emissions:output ratio  $\phi_t$  and chosen emissions control level  $\tau_t$ .

$$E_t = (1 - \tau_t)\phi_t Y_t \tag{12}$$

 $<sup>^{6}</sup>$ For a discussion of the effects of utility function choice on the outcome of learning dynamics, see Gollier et al. (2000)



Figure 3: Schematic Diagram of the Climate Model

The emissions: output ratio  $\phi$  is assumed to converge over time, and is characterized by the law of motion:

$$\phi_t = \phi_0 \, \exp\left(\frac{\gamma_\phi}{\delta_\phi} (1 - e^{-\delta_\phi t})\right). \tag{13}$$

Carbon concentration evolves as a function of the existing stock and period emissions. The law of motion is operationalized in (14) with a decay rate ( $\delta_m \in [0, 1)$ ) and a preindustrial baseline carbon concentration  $m_b$ .

$$m_t = E_t + (1 - \delta_m)(m_t - m_b) + m_b.$$
(14)

 $CO_2$  stocks lead to a change in radiative forcing, increasing heat retention relative to the baseline for the atmosphere. Radiative forcing,  $F_t$ , is modelled as a function of atmospheric carbon as follows:

$$F_t = \eta \left[ \frac{\log \frac{m_t}{m_b}}{\log(2)} \right]. \tag{15}$$

where the units are the relative increase in radiative forcing from pre-industrial times in  $\frac{W}{m^2}$ . The parameter  $\eta$  has units of  ${}^{o}C$  and represents the long-run surface temperature increase from a doubling of atmospheric CO<sub>2</sub> relative to pre-industrial levels. Global surface and ocean temperature changes are influences by the amount of radiative forcing, and current temperatures as defined in (1-2).

#### 4.3 Benchmark Model

The social planner's problem is to choose a sequence of savings and emissions control rates to maximize the net present value of aggregate utility. The maximization is budget constrained, such that each period's consumption and investment cannot jointly exceed production. Uncertainty in the benchmark economy occurs in the stochastic, autoregressive process for temperature change. The value function is allowed to be time-varying for T-1 years, after which it is assumed to be constant, such that T separate value functions define the recursive problem. The Social Planner's recursive problem is outlined below<sup>7</sup>:

$$V_t(K, G, O, m) = \max_{\psi, \tau} L(\frac{\frac{C}{L}}{1 - \sigma})^{1 - \sigma} + \beta E \left[ V_{t+1}(K', G', O', m') \right]$$
(16)

s.t. 
$$C = \psi * \frac{1 - b_1 \tau^{b_2}}{1 + \theta_1 G^{\theta_2}} A K^{\alpha} L^{1 - \alpha}$$
 (17)

$$K' = (1 - \delta_k)K + (1 - \psi) * \frac{1 - b_1 \tau^{b_2}}{1 + \theta_1 G^{\theta_2}} A K^{\alpha} L^{1 - \alpha}$$
(18)

$$m' = (1 - \tau)\theta A K^{\alpha} L^{1 - \alpha} + (1 - \delta_m)(m - m_b) + m_b$$
(19)

$$G' = \lambda_1 G + \eta \left[\frac{\log \frac{m_t}{m_b}}{\log(2)}\right] + \omega O + u, \ u \sim N(0, \sigma_u^2) \tag{20}$$

$$O' = \lambda_2 O + (1 - \lambda_2)G \tag{21}$$

$$V_T(K, G, O, m) = \max_{\psi, \tau} L\left(\frac{\underline{C}}{1-\sigma}\right)^{1-\sigma} + \beta E\left[V_T\left(K', G', O', m'\right)\right]$$
(22)

where  $\beta$  gives the discount factor  $\beta \in (0, 1)$ .

#### 4.4 Learning Model

With the learning characterization embedded in the IAM, the planner is learning over two parameters,  $\lambda_1$  and  $\eta$ , or the autoregressive and radiative forcing components of temperature change. Prior means are denoted by  $\hat{\lambda}_1$  and  $\hat{\eta}$ , and the prior variance-covariance matrix by **V**. In order to learn about two parameters and their variances, the planner must be able to observe current temperatures and climate variables as well as lagged values of each for three periods. The state space of the planners problem expands to  $S = \{K, G, G_{-1}, G_{-2}, O, O_{-1}, O_{-2}, m, m_{-1}, m_{-2}, \hat{\lambda}_1, \hat{\eta}, \mathbf{V}\}$ , and denoting:

$$\begin{array}{ccc} G & \ln \frac{m}{m_b} & G' - \omega O_t \\ X = \begin{pmatrix} G_{-1} & \ln \frac{m_{-1}}{m_b} \\ G_{-2} & \ln \frac{m_{-2}}{m_b} \end{pmatrix}, \quad H = \begin{pmatrix} G - \omega O_{t-1} \\ G_{-1} - \omega O_{t-2} \end{pmatrix}$$

the recursive problem can be written as:

<sup>&</sup>lt;sup>7</sup>Time subscripts have been removed from the state and control variables. For example, for capital, K indicates today's capital stock, K' is next period's stock and  $K_{-1}$  is capital stock with a one period lag. L, A, and  $\phi$  are functions of the current calendar time t as given in (7), (8), and (13)

$$V_t(S) = \max_{\psi,\tau} L(\frac{\frac{C}{L}}{1-\sigma})^{1-\sigma} + \beta E\left[V_{t+1}(S')\right]$$
(23)

s.t. 
$$C = \psi * \frac{1 - b_1 \tau^{\theta_2}}{1 + \theta_1 G^{\theta_2}} A K^{\alpha} L^{1-\alpha}$$
 (24)

$$K' = (1 - \delta_k)K + (1 - \psi) * \frac{1 - b_1 \tau^{b_2}}{1 + \theta_1 G^{\theta_2}} A K^{\alpha} L^{1 - \alpha}$$
(25)

$$m' = (1 - \tau)\theta A K^{\alpha} L^{1-\alpha} + (1 - \delta_m)(m - m_b) + m_b$$

$$\log \frac{m_t}{m_b}$$
(26)

$$G' = \lambda_1 G + \eta \left[\frac{\log \overline{m_b}}{\log(2)}\right] + \omega O + u, \ u \sim N(0, \sigma_u^2) \tag{27}$$

$$O' = \lambda_2 O + (1 - \lambda_2)G \tag{28}$$

$$G'_{-1} = G \tag{29}$$

$$G'_{-2} = G_{-1} \tag{30}$$

$$\begin{pmatrix} \hat{\lambda}'_{1} \\ \hat{\eta}' \end{pmatrix} = \left[ \mathbf{V}^{-1} + \left[ \sigma_{u}^{2}(X^{T}X) \right]^{-1} \right]^{-1} \left[ \mathbf{V}^{-1} \begin{pmatrix} \hat{\lambda}_{1} \\ \hat{\eta} \end{pmatrix} + \left[ \sigma_{u}^{2}(X^{T}X)^{-1} \right]^{-1} (X^{T}X)^{-1} X^{T} H \right]$$
(31)

$$\mathbf{V}^* = \begin{bmatrix} \mathbf{V}^{-1} + \begin{bmatrix} \sigma_u^2 (X^T X)^{-1} \end{bmatrix}^{-1} \end{bmatrix}^{-1}$$
(32)

$$V_T(S) = \max_{\psi,\tau} L(\frac{\overleftarrow{L}}{1-\sigma})^{1-\sigma} + \beta E \left[ V_T(S') \right]$$
(33)

The parameters of the model are  $\{b_1, b_2, \theta_1, \theta_2, \lambda, \delta_k, \delta_m, \delta_a, \delta_l, \delta_\phi, \gamma_a, \gamma_l, \gamma_\phi, m_b, \lambda_1, \eta, \omega, \sigma_u^2\}$ . There still must exist a true  $\lambda_1$  and  $\eta$  to simulate the path of the learning economy even though the planner's decisions are based on the parameter estimates of  $\hat{\lambda}_1$  and  $\hat{\eta}$ .

#### 4.5 Equilibrium and Dynamics

**Definition 4.5.1** Dynamic equilibrium in the model is a sequence of beliefs  $\hat{\lambda}_1$  and  $\hat{\eta}^8$ , control, state and climate variables such that given beliefs about the climate change mechanism, the social planner maximizes expected sum of per capita utility through the sequence of choices of emissions control and savings rates. For both the benchmark and learning models, this equilibrium is characterized by the fixed point solution to Bellman's equation in each time period.

# 5 Computation

The model presented above is a stylized representation of the interaction between a socially planned population of agents and their physical environment. In attempting to account for the dynamics of climate, economy and knowledge, the minimum number of state variables in the recursive problem

<sup>&</sup>lt;sup>8</sup>In the benchmark economy, the beliefs are the true values in each time period

is 10, leading to a continuous state space defined over  $\mathbb{R}^{10}$ . In order to solve the value function in (23), traditional value function and policy function iteration approaches will not provide accurate estimates due to the coarseness of the state space grid that would be required. This is referred to as the curse of dimensionality, where the grid on which the solution is computed grows exponentially with the dimension of the state space for a given level of accuracy.

Pizer (1996) proposes the use of a linear-quadratic approximation, however this method is only valid in the neighborhood of the steady state of the economy, and thus does not provide a good estimate of the value function over the entire scope of time in question.

Kelly and Kolstad (1999b) address this curse of dimensionality by using a neural networks approach to approximate the value function between grid points. While the neural network approach reduces the curse of dimensionality, it continues to involve the solution to a non-linear minimization problem.

This paper makes use of randomization combined with a regression-based interpolation algorithm analogous to those proposed in Keane and Wolpin (1994) and Reiter (1999). The full description of this algorithm is presented in Leach (2003b), but the outline is provided below:

#### Algorithm 1

Objective: Compute an approximation to the value function using a recursive, iteration algorithm and approximation by regression.

Algorithm Preliminaries: Choose a convergence criterion  $\epsilon$ , a time horizon T and an initial guess of the value function V0. Generate a discretized grid over the values of the state variables. The value function will be defined over this grid. ( $\mathbf{K} \times \mathbf{ST} \times \mathbf{OT} \times \mathbf{mAT}$ ). Choose a transformation function for the state variables  $g(\mathbf{S})$  to be used as an approximating regression.

Step 1: Using the function  $g(\mathbf{S})$  and the values at the grid points, calculate approximation coefficients using OLS.

Step 2: For each point on the grid, solve the maximization problem given in (23). Use derivatives of the return function and  $\frac{\partial g(\mathbf{S})}{\partial K'}$  and  $\frac{\partial g(\mathbf{S})}{\partial mAT'}$  to generate a system of two zero functions in the two unknowns  $\psi_t$  and  $\tau_t$ . Solve for  $\psi_t$  and  $\tau_t$  taking account of potential corner solutions. (see above) Formally, update the value at each point on the grid as:

$$v_t^{n+1}(K, ST, OT, mAT) = \max_{K' \in \mathbf{K}} \left\{ U(K') + \beta Eg(K', K, ST, OT, mAT) \right\}$$
(34)

Step 2: Continue iterating until the difference between  $v^n$  and  $v^{n+1}$  satisfies the convergence criteria.

The algorithm used in this paper relies on the OLS estimators as the global approximation of the value function. In order to evaluate the integrals in the planners problem in both the benchmark and learning cases, Monte Carlo methods are used. In order to restrict the number of state variables and keep the problem tractable (10 states), an additional assumption is made on the learning dynamics in order to compute the value function. Agents are assumed not to keep track of the lagged values of either ocean temperature or atmospheric carbon. Instead, when they enter a new period and observe the evolved values of temperature, carbon, and ocean temperature, it is assumed that they calculate growth rates for surface temperature ( $\gamma_G$ ), ocean temperature ( $\gamma_O$ ) and atmospheric carbon forcing ( $\gamma_f$ ) and use those growth rates to calculate the time path for forcing and ocean temperature for the previous two years which they use to update their beliefs conditional on observed temperature and lags.

$$X = \begin{pmatrix} G & \ln \frac{m}{m_b} \\ G_{-1} & \frac{\ln \frac{m}{m_b}}{(1+\gamma_f)} \\ G_{-1}(1+\gamma_G)^{-1} & \frac{\ln \frac{m}{m_b}}{(1+\gamma_f)^2} \end{pmatrix}, \quad H = \begin{pmatrix} G' - \omega O_t \\ G - \omega \frac{O_t}{(1+\gamma_O)} \\ G_{-1} - \omega \frac{O_t}{(1+\gamma_O)^2} \end{pmatrix}$$

This assumption allows for the elimination of 5 state variables from the planner's problem, which renders the problem more tractable with minimal approximation error. The sensitivity of the learning path to this assumption was tested using Monte Carlo methods. Learning paths were simulated for the full information case and the assumption stated above. From Figure 4, it is evident that the time paths are very similar and the bias introduced by the assumption is decreasing and positive, with a maximum of 3% which leads to a small increase in learning times, which is consistent with a small reduction in information used. In Figure 5, the magnitude of the bias is shown to be decreasing through time. These assumptions hold little cost in terms of accuracy because temperature, ocean temperature and atmospheric carbon are all stocks with very small rates of change, and the second derivatives of these stocks are very close to zero for all time periods. It therefore seems a reasonable assumption to set the second derivatives to zero and assume that the rates of change are constant in order to reduce by one third the number of state variables in the model.

# 6 Calibration and Simulations

The model is calibrated using parameters taken from a combination of the model presented in Leach (2003a) and Pizer (1996), each of which feature an annual time interval and are thus directly comparable to the model contained in this paper. The calibration is the same for the benchmark and learning model, except that the prior mean values are the assumed true values in the benchmark



Figure 4: Learning Paths with Full and Imputed Information

model, and there is virtual certainty about these parameters. The economy is calibrated to a global production function as outlined in Pizer(1996). Climate sector calibrations are derived from Pizer and Kelly and Kolstad (1999a), and updated to take account of the annual rather than 10 year time interval. Variance of temperature residuals was taken from an estimation of the proposed climate model using temperature and carbon stock data from the CDIAC. Simulation starting values are shown in Table 1, and parameter calibrations are in Table 2.

#### 6.1 Learning Times with Endogenous Emissions

In order to quantify the role of learning in the optimal regulation of climate change, the model is simulated for three scenarios. The first scenario is the benchmark, where the social planner knows with certainty the parameters which govern the law of motion of surface temperature, and thus only



Figure 5: Approximation Error for Learning Paths

forms expectations over the random shock to temperature in each period.<sup>9</sup> The second scenario allows learning over the forcing parameter only. The prior mean for the autoregressive parameter has the correct value and a variance of  $10^{-6}$ . The planner's estimate of the forcing parameter admits an error and uncertainty. The third scenario shows the effect of uncertainty over both the severity of climate sensitivity to atmospheric carbon levels and the persistence of shocks. These scenarios allow for the identification of exactly the type of effect discussed in Gollier et al. (2000): the precautionary motive associated with risk over future damages. Each of these is discussed in turn below.

Climate evolution in this model is governed by two key parameters: the long run sensitivity of surface temperature to a doubling in atmospheric carbon and the coefficient of autoregression on

 $<sup>^{9}</sup>$ This paper makes no claim about the true law of motion. In the simulations, there is a true parameter which governs the evolution of surface temperature, and this is known to the regulator.

surface temperature deviations. For the benchmark scenario, the regulator knows with certainty that these parameters are  $3^{\circ}C$  and .9112 respectively.<sup>10</sup>

In the second scenario, learning over a single parameter is examined. The regulator has an error of  $\pm 1^{\circ}C$  for the climate sensitivity parameter and correct prior mean for the autoregressive parameter. The prior variance leads to 99% confidence intervals of  $\pm 2.5^{\circ}C$  and .00001 on the two parameters respectively. The probability density function for the forcing parameter corresponds to those published in Knutti et al. (2002), and also to confidence intervals estimated using CDIAC temperature and carbon concentration data. Prior covariance is set to 0. This scenario is designed to capture the role of parametric uncertainty around a single parameter on optimal policy choices.

In the third scenario, the persistence of an incorrect estimation of the climate sensitivity to carbon levels is examined when two parameters are uncertain. Relative to assumed true parameters, this scenario begins again with an error of  $\pm 1^{\circ}C$  on the parameter governing the temperature change from a doubling of CO<sub>2</sub> in an effort to compare how long this error is persistent when combined with uncertainty in the autoregressive component. The confidence levels are wide enough such that the 99% bands admit  $\pm 2.5^{\circ}C$  and  $\pm 0.025$  values for each of the parameters. This simulation is aimed at evaluating two characteristics of the problem: the time to eliminate an error of a given magnitude and the effect of erroneous estimates and uncertainty on optimal policy.

Figures 6 and 7 show the evolution of the prior mean estimates for parameters governing the autoregressive process and radiative forcing respectively.

In the first scenario, we can see that the prior mean value converges immediately toward the true value for the forcing parameter. The time to learn the true value is defined here by the time it takes to reject a value of  $.5^{\circ}C$  away from the true climate sensitivity value at the 99% confidence level. In this case, it takes over 250 years for learning to take place. In terms of chosen policy, the differences are positive but marginal. In the first scenario, a 1% positive (negative) squared deviation in the estimated forcing parameter from the true forcing parameter is expected to bring about a 4.3% increase (decrease) in regulation squared. This estimate has a standard error of .2115.

In terms of time to learn in the second scenario, the initial prior mean for the forcing parameter differs by  $1^{\circ}C$  from the true parameter. In the simulation, it takes over 1000 years for the planner to reject a value of  $.5^{\circ}C$  away from the true parameter. These stark increases in learning times over the exogenous emissions growth rates case occurs for two reasons. First, the planner does not directly account for learning in the emissions control decision, so by controlling emissions at a greater rate

<sup>&</sup>lt;sup>10</sup>Certainty here implies a prior variance of  $10^{-5}$ .



Figure 6: Prior mean - AR(1) parameter  $\hat{\lambda}_1$ 

as stocks tend to increase, the regulator is endogenously lengthening learning times. Second, the convergence assumptions on the economy lead to emissions growth rates with are smaller than 0.2% for most of the learning simulation.

The third driving force behind this divergence in learning times is the self-fulfilling hypothesis. We can see how an initial over-estimation of the forcing effect leads us to immediately adjust our estimate of the AR(1) parameter downward. We continue to update in this way, but the Bayesian learning structure is not able to sufficiently separately identify the effect of carbon from the autoregressive effect, to the forcing parameter will converge initially and then stabilize above it's true value, while the AR(1) parameter will diverge from it's true value initially and then stabilize below it's true value. These errors will be persistent as long as there continues to be a stock of carbon in the atmosphere. Once this stock has decayed, there is only one potential effect (AR(1)) and the learning mechanism is able to identify its role. It is possible for learning about the climate



Figure 7: Prior mean - Radiative forcing parameter  $\hat{\eta}$ 

change parameter to occur such that the planner is eventually able to exclude the true parameter.

The differences in chosen policy are larger here than in the first scenario. It is predicted that a 1% positive (negative) squared deviation in the estimated forcing parameter from the true forcing parameter is expected to bring about a 6.3693% increase (decrease) in regulation squared. This estimate has a standard error of .2205, so it is significant at any standard level.

Figure 8 shows the absolute differences in optimal control rates through time, based on the benchmark simulation where the regulator is certain about the true parameters of climate change. In both scenarios, the differences in chosen policy are small, as the optimal level of abatement changes only marginally in  $\eta$ . Policy choice is positively related to the mean estimate  $\hat{\eta}$ . Clearly, the deviations in optimal policy are small, and are dependent on the mean of the prior distribution. This result is similar to that found in Kolstad and Kelly (1999) although the persistence in mean differences is higher due to the learning over two parameters.



Figure 8: Differences in Optimal Control-Benchmark and Scenarios

#### 6.2 Regulation with a certainty rule

# 7 Conclusion

This paper is an exercise aimed at exploring two questions. First, how does learning about climate change vary when there exists uncertainty over more than one parameter. Second, this paper is able to compare the optimal policy sequences in models with and without uncertainty and learning to examine the quantitative importance of uncertainty in this dimension. Tangentially, this paper also implements a new algorithm for solving integrated assessment models under uncertainty which makes use of the recursive structure of the problem and is less susceptible to the curse of dimensionality.

In the first exercise, it is abundantly clear that the time to learn the true value of a parameter increases exponentially with the number of parameters being investigated, and with the number of alternatives the agent can consider. In Kelly and Kolstad (1999a) when the learning involved the consideration of two opposing hypotheses about a single parameter, the learning time was shown to be on the order of 1-25 periods. When the assumption of a discrete distribution is removed, and learning occurs about a single parameter with a continuous distribution, the time period increases to 50-200 years. Further examination of learning about two continuous parameters shows that the time to learn can be on the order of thousands of periods.

In response to the second question, little difference is shown between optimal policy sequences when learning is admitted to the model as discussed. Furthermore, it is shown that, in a simulated growth model, learning times are much higher when considering two parameters with uncertain values. In the model proposed herein, it takes greater than 250 years for an error of  $2^{0}C$  in the climate sensitivity to carbon to be reduced to  $.1^{0}C$ , however the convergence of a policy rule may be substantially faster. In a simulation in which the benefits of a policy to learn-then-act are compared to one of act-then-learn, the former is shown to dominate.

Much of the literature focusses on the question of whether to act-then-learn, or learn-thenact. This model is not able to provide a positive answer to this question, since the losses to either strategy are conditional on the initial prior error. What is important to note however is the decrease in learning times which arise from the learn-then-act strategy. In this case, the emissions control policy does not provide a dampening effect on the variance of emissions, thus learning times are decreased.

The end results of this paper are best summarized by emphasizing the exponential response of learning times to uncertainty. Climate change is a process over which there exists great uncertainty, and it is of paramount importance to treat it as completely as possible. With previous characterizations of learning and climate change, the quoted learning times can be interpreted that there may be some reward to maintaining high emissions growth rates, since this will allow us to create optimal policy. When uncertainty is modelled in a more complete treatment, this motive quickly disappears.

## References

CYERT, R. M., AND M. H. DEGROOT (1974): "Rational Expectations and Bayesian Analysis," Journal of Political Economy, 82, 521–535.

DOORNIK, J. A. (2003): Ox Version 3.30. Oxford.

- EPSTEIN, L. (1980): "Decision Making and the Temporal Resolution of Uncertainty," International Economic Review, 21, 269–83.
- KEANE, M., AND K. WOLPIN (1994): "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," *Review of Economics and Statistics*, 76(4), 648–672.
- KELLY, D. L., AND C. D. KOLSTAD (1999a): "Bayesian Learning, Growth and Pollution," Journal of Economic Dynamics and Control, 23, 491–518.

(1999b): "Solving Growth Models with an Environmental Sector," Forthcoming, Computational Economics.

- KELLY, D. L., C. D. KOLSTAD, AND G. MITCHELL (1999): "Adjustment costs from Environmental Change Induced by Incomplete Information and Learning," Working Paper.
- KNUTTI, R., T. F. STOCKER, F. JOOS, AND G.-K. PLATTNER (2002): "Constraints on radiative forcing and future climate change from observations and climate model ensembles," *Nature*, 416, 719–723.
- KOLSTAD, C. D. (1996): "Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions," Journal of Environmental Economics and Management, 18, 1–18.
  (1997): The Economics of Global Warming chap. The Timing of CO<sub>2</sub> Control in the
- Face of Uncertainty and Learning, pp. 523–544, The International Library of Critical Writings in Economics. Elgar Reference Collection, Cheltenham, UK.
- KONDRATYEV, K. Y. (1988): *Climate Shocks: Natural and Anthropogenic*, Wiley Series in Climate and the Biosphere. John Wiley and Sons, New York, USA.
- MADDISON, D. (1995): "A cost-benefit analysis of slowing climate change," *Energy Policy*, 23, 337–46.
- MANNE, A. S., R. MENDELSOHN, AND R. RICHELS (1995): "MERGE: A model for evaluating regional and global effects of GHG reduction policies," *Energy Policy*, 23(1), 17–34.
- MANNE, A. S., AND R. RICHELS (1992): Buying Greenhouse Insurance: The Economic Costs of CO<sub>2</sub> Emissions Limits. The MIT Press, Cambridge, Mass.
- NORDHAUS, W. D. (1994): Managing the Global Commons: The Economics of Climate Change. MIT Press, Cambridge, Mass.
- NORDHAUS, W. D., AND J. BOYER (2000): Warming the World. MIT Press, Cambridge, Mass.
- PIZER, W. A. (1996): "Modeling Long-Term Policy Under Uncertainty," Ph.D. Thesis Harvard University.
- RAMSEY, F. (1928): "A Mathematical Theory of Saving," The Economic Journal.

- REITER, M. (1999): "Solving higher-dimensional continuous time stochastic control problems by value function regression," *Journal of Economic Dynamics and Control*, 23, 1329–53.
- RIOS-RULL, J. V. (1997): "Computation of Equilibrium in Heterogeneous Agent Models," Working Paper - University of Pennsylvania.
- ULPH, A., AND D. ULPH (1997): "Global Warming, Irreversibility and Learning," *The Economic Journal*, 107, 636–650.

Table 1: State Variables Used in the Model				
State	Definition	Units	Starting	
			Value(BM)	
K	Capital Stock	$10^{1}2$ \$1987	12	
M	Atmospheric $CO_2$	$10^9$ tonnes	672.49	
G	Global Surface Temperature	$^{o}C$ from mean	.184	
0	Ocean Temperature	$^{o}C$ from mean	.02	
$T_{-1}$	Temperature Lag 1	$^{o}C$ from mean	.146	
$\mu_1$	Mean Estimate $\hat{\alpha}$	-	.9112	
$\mu_2$	Mean Estimate $\hat{\beta}$	$^{o}C/w/m^{2}$	3	
$\mathbf{V}_{11}$	Variance Estimate $\hat{\alpha}$	-	0	
$\mathbf{V}_{22}$	Variance Estimate $\hat{\beta}$	$(^{o}C/w/m^{2})^{2}$	0	
$\mathbf{V}_{21,12}$	Covariance Estimate	$^{o}C/w/m^{2}$	0	

Table 2: Calibrated Values

D	Table 2: Calibrated Values	
Parameter	Description	Calibrated Value
$\sigma$	Coefficient of Relative Risk Aversion	.95
eta	Discount Rate	.95
δ	Capital Depreciation Rate	.1
$\alpha$	Production Share of Capital	.3764
$\gamma_a$	Growth rate of labour augmenting tech. change	.008
$\delta_a$	Decay rate on $\gamma_a$	.011
$\gamma_n$	Growth rate of population	0.0203
$\delta_n$	Decay rate on $\gamma_n$	.017
$\gamma_{\phi}$	Growth rate of emissions:output ratio	-0.0125
$\delta_{\phi}$	Decay rate of emissions:output ratio	.011
$m_b$	Preindustrial concentration of $CO_2$ (GtC)	590
$(1 - \delta_m)$	atmpspheric retention of $CO_2$	.99167
$\lambda_1$	Autoregressive component of temp. change	.9112
$\lambda_2$	AR(1) parameter in ocean temperature	.9912
$\eta$	Temperature sensitivity to $CO_2$ doubling ( $^{o}C$ )	3
$b_1$	Linear control costs	.0686
$b_2$	Exponential control costs	2.877
$\sigma_u$	Standard Deviation of Temp. Residual	.1233
ω	Coefficient on Ocean Temp in Temperature Change	.0002
$ heta_1$	Linear Component of Damages	.001478
$ heta_2$	Exponent in Damage Function	2

# Liste des cahiers de recherche publiés par les professeurs des H.E.C. 2003-2004

# Institut d'économie appliquée

- IEA-03-01 ROBERT GAGNÉ; PIERRE THOMAS LÉGER. « Determinants of Physicians' Decisions to Specialize »,29 pages.
- IEA-03-02 BENOIT DOSTIE. « Controlling for Demand Side Factors and Job Matching: Maximum Likelihood Estimates of the Returns to Seniority Using Matched Employer-Employee Data »,24 pages.
- IEA-03-03 ALAIN LAPOINTE. « La performance de Montréal et l'économie du savoir: un changement de politique s'impose », 35 pages.
- IEA-03-04 MICHEL NORMANDIN; LOUIS PHANEUF. « Monetary Policy Shocks: Testing Identification Conditions Under Time-Varying Conditional Volatility », 43 pages.
- IEA-03-05 MARTIN BOILEAU; MICHEL NORMANDIN. « Dynamics of the Current Account and Interest Differentials », 38 pages.
- IEA-03-06: MICHEL NORMANDIN; PASCAL ST-AMOUR. « Recursive Measures of Total Wealth and Portfolio Return », 10 pages.
- IEA-03-07: DOSTIE, BENOIT ; LÉGER, PIERRE THOMAS. « The Living Arrangement Dynamics of Sick, Elderly Individuals », 29 pages.
- IEA-03-08: MICHEL NORMANDIN. « Canadian and U.S. Financial Markets: Testing the International Integration Hypothesis under Time-Varying Conditional Volatility », 35 pages.

- IEA-04-01: ANDREW LEACH. « Integrated Assessment of Climate Change Using an OLG Model », 34 pages.
- IEA-04-02: ANDREW LEACH. « SubGame, set and match. Identifying Incentive Response in a Tournament », 39 pages.