Les Cahiers du GERAD ISSN: 0711–2440

Limited Predictability of a THC Shutdown: Implications for Economic Policy

O. Bahn
A. Leach

G–2005–46

May 2005

Limited Predictability of a THC Shutdown: Implications for Economic Policy

Olivier Bahn
GERAD and Méthodes quantitatives de gestion
HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada, H3T 2A7
olivier.bahn@hec.ca

Andrew Leach
CRT, CIRANO, CIRPEE
and Institut d'économie appliquée
HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada, H3T 2A7
andrew.leach@hec.ca

May 2005

Les Cahiers du GERAD

G–2005–46

Abstract

We modify an existing model of climate and economy to address the effect of uncertain, threshold events on the choice of optimal emissions control policy. We augment an existing model to include a non-linear response to climate system perturbations modeled on the potential shutdown of the Atlantic thermohaline circulation (THC). Using a model which features decision making under uncertainty by a social planner, we are able to quantify the cost of parameter uncertainty and judge the value of perfect information. We find that uncertainty over climate sensitivity to carbon emissions and over the value of the threshold has minimal costs given standard economic growth model parameters, which is counterintuitive given assumptions and conclusions in previous papers. Additionally, we find that important changes occur as a result of removing the assumption that non-linearities in damages occur in tandem with non-linearities in the climate system.

Key Words: Climate Change Environmental Regulation; Growth; Pollution; Dynamic Programming; Abrupt Climate Changes; Thermohaline Circulation.

Résumé

Nous modifions un modèle existant du climat et de l'économie pour étudier les effets d'évènements à seuil incertain sur le choix de la politique optimale de contrôle des émissions. Plus précisément, nous améliorons un modèle existant pour y inclure une réponse non linéaire aux perturbations du système climatique modélisée sur une rupture potentielle de la circulation thermohaline. À l'aide d'un modèle qui comporte la prise de décision sous incertitude par un planificateur social, nous pouvons quantifier le coût de l'incertitude et évaluer la valeur de l'information parfaite. Nous trouvons que l'incertitude liée à la sensibilité du climat aux émissions de carbone et à la valeur du seuil à un coût très faible étant donnés les paramètres standards du modèle de croissance économique, ce qui contredit les hypothèses et conclusions de papiers précédents. En outre, nous trouvons que d'importants changements se déroulent lorsque l'on supprime l'hypothèse que les non linéarités dans les dommages vont de pair avec les non linéarités dans le système climatique.

Acknowledgments: The authors acknowledge support from the Direction de la Recherche, HEC Montréal. Thanks to seminar participants at the HEC Montréal Summer Research Workshop and UQAM. Olivier Bahn acknowledges the support of NSERC.

1 Introduction

We modify an existing model of climate and economy to address the effect of uncertain, threshold events on the choice of optimal emissions control policy. An emerging literature on the economics of climate change has begun to address the effect of events which occur with limited predictability and which bring about long-term if not permanent changes in our climate system. To this end, the literature has taken two distinct approaches, characterizing emissions policies which will prevent such an irreversible event from occurring, or examining the implications of foreseeable extreme events on optimal policy choices. We first attempt to reconcile these two approaches by evaluating the costs to inaction and overreaction in the face of threshold levels in the climate system. Second, we evaluate the value of information in a climate system with thresholds. The model we propose differs from others in the literature, in that non-linearities in the climate system are not directly related to non-linearities in damages. Rather, we assume that damages are driven by the state of world climate, not the threshold itself.

We focus on potential shutdown thresholds in the Atlantic thermohaline circulation (THC), a global scale ocean circulation pattern and an important component of the climate system. It is mainly driven by density differences in sea water, and as such by temperature and salinity (Wunsch, 2002). Several studies indicate that THC may be strongly reduced in response to global warming (see Cubasch and Meehl, 2001, for a review). Such a collapse in the THC would yield a radically different climate in the North Atlantic region (see e.g. Vellinga and Wood, 2002) with potentially severe economic and social consequences (see e.g. National Research Council, 2002).

A potential THC shutdown is generally characterized in the literature as the result of the global climate system crossing some threshold, which triggers a complete and inevitable shutdown. A key characteristic of this shutdown is that it occurs over several hundred years (e.g., see Manabe and Stouffer, 1993; Schmittner and Stocker, 1999). Recovery of the THC from a near-collapse also occurs over a very long time horizon. Furthermore, the nature of the threshold is uncertain and multi-dimensional (e.g., see Knutti and Stocker, 2002). Indeed, regardless of temperature level, a short period of rapid temperature increase could trigger a shutdown, while the same change in temperature realized over a longer period would not (Stocker and Schmittner, 1997.

In this paper, we present an integrated assessment model (IAM) which takes account of the particular characteristics of a THC shutdown. IAMs typically combine key elements of the economic and biophysical systems, elements that underlie the anthropogenic global climate change phenomenon. Several studies conducted with IAMs, such as DICE (Nordhaus, 1994) and MERGE (Manne et al., 1995), consider the generic possibility of abrupt climate changes, but not irreversible ones. In the infinite horizon analog to the Nordhaus and Boyer (2000) model in particular, any changes to the climate are completely reversible, as temperature is modeled as a deterministic, autoregressive process which can be influenced by greenhouse gas (GHG) emission levels. As soon as emissions decline to zero, temperature will revert to its pre-industrial average. The science of climate change is often at odds with this characterization, arguing that certain effects of climate change may be irreversible, or at least irreversible over planning horizons many orders of magnitude longer than those relevant to economic models. The THC shutdown problem provides an excellent example of just such a case.

A few papers have explicitly taken into account a possible, irreversible collapse of the THC. Specifically, we highlight below our contribution in light of Keller et al. (2000, 2004), Mastrandrea and Schneider (2001), and Zickfeld and Bruckner (2003). We discuss briefly the approaches taken in these papers with respect to the links between climate and economic systems and their treatment of uncertainty.

In terms of tracking the effect of economic activity on climate variables, Mastrandrea and Schneider (2001) and Zickfeld and Bruckner (2003) use climate models to model changes in the THC in response to forcing agents (greenhouse gases). In Keller et al. (2000, 2004), the THC shuts down when the climate system crosses a calibrated threshold, and damages increase as a result. While our paper does not employ a highly detailed model of global climate, we include laws of motion which capture the evolution of the THC in response to climate forcing agents. These laws of motion are calibrated to experiments performed using a coupled ocean-atmosphere circulation model, Bern 2.5-D (Stocker et al., 1992) in Stocker and Schmittner (1997).

Policy analysis for greenhouse gas emissions control will depend very strongly on the characterization of the relationship between the state of world climate and the productivity of the world's economy. In Zickfield and Bruckner (2003), damages from a THC shutdown are implicitly assumed to be prohibitive, since they seek to characterize policies preventing a THC collapse. In Keller et al. (2000,2004) and Mastrandrea and Schneider (2001), climate changes alter the productivity of the economy through a damage function, allowing these papers to characterize the costs and benefits of emissions control policies. Keller et al. (2000, 2004) assume that damages occur as a result of crossing the threshold, not as a function of the remaining circulation. They modify the original Nordhaus (1994) model with an additional parameter which reduces productivity in the economy in all periods after the threshold has been crossed. Mastrandrea and Schneider (2001) argues that consideration of the effects of the THC shutdown should lead to a damage function which increases more rapidly in temperature change than that used in the original Nordhaus (1994) paper. Their analysis thus summarizes the relationship between the equilibrium paths of temperature, THC, and damages through a single damage parameter. The additional damages characterized by Mastrandrea and Schneider (2001) are valid along the transition path of the converged, optimal policy model, but are not necessarily consistent with damages off this path. As a result of our direct mapping between predicted circulation and damages (as well as damages from temperature change), we are able to perform policy experiments away from the optimal transition path of the model. Our characterization differs also from

Keller et al. (2000, 2004) in that we are able to capture damages which would be present under a near-collapse of the THC, as damages are stipulated as a function of the loss of circulation, not on the crossing of the threshold.

The importance of the above should not be understated. In each of the models listed above, a planner is assumed to care about agents' utility from consumption, which is negatively affected by climate changes. Using a model in which crossing a climate threshold has an immediate and drastic effect on the production possibility frontier of the economy should be considered an important assumption. The scientific literature is by no means clear on the question of whether non-linearities in damages coincide directly with the crossing of thresholds in the climate system. With this in mind, we proceed with a model where the threshold level of the THC affects only future possible values of the THC, while damages are related not to the threshold itself, but to the remaining fraction of circulation. Imposing a direct consumption cost of crossing a threshold may bias the planner in favour of the choice of THC-preserving policies. In our results, we find negative welfare effects of a THC-preserving policy. This allows our results to add greater context to papers which examine the constraints that must be placed on the economy in order to preserve the THC. Furthermore, the fact that the planner is willing to accept a THC collapse in our simulations with a 2.5% reduction in productivity as compared to a critical value of less than 1% above which Keller et al. (2000) suggests THC-preserving policies to be optimal captures the importance of assumptions on the damage function.

The previous literature has led to the belief that the effect of uncertainty will be very important in the selection of optimal policy in the face of extreme events, particularly when an irreversible threshold triggers such an event. Keller et al. (2004) reports on a stochastic programming approach that considers three different settings for uncertain parameters that define different possible future states of the world to which subjective probabilities are associated. In this environment, the optimal policy is chosen conditional on uncertainty, but this uncertainty is resolved completely in the future. This model shows significant differences in optimal policy choices conditional on uncertainty. This method evaluates an optimal hedging policy against a known set of future possibilities, however the stochastic programming approach limits the choice of uncertainty to discrete probabilities over a small number of scenarios.

Casual empiricism suggests that climatologists are more at ease with providing probability density functions for the uncertain parameters than probabilities for a few contrasted scenarios. Keller et al. (2004) also reports on a probabilistic optimization procedure which addresses uncertainty on two key parameters using a Monte Carlo simulation. The optimal policy solution here is that which maximizes the expected indirect utility. This an approximation to the solution we will examine, as we will directly evaluate the indirect utility of each state of the world given uncertainty. While stochastic programming and probabilistic optimization each provide information on the effect of uncertainty, they do not specifically take account of the effects of the precautionary motive which may be present in the social planner's behaviour over the entire time-span of interest. In order to solve the model under uncertainty, we characterize the value function of the social planner, and solve iteratively for this function using techniques developed in Kelly and Kolstad (2001). In such an environment, we find that the indirect utility function for the social planner at relevant states along the transition path is largely linear in future values of the THC. As such, the costs of uncertainty around the transitions of this variable are small.

The remainder of this paper proceeds as follows. Section 2 proposes the model of climate and economy. Calibrations are discussed in Section 3. The model is solved in Section 4 and simulated for three uncertainty scenarios in Section 5, and Section 6 concludes.

2 The Model

The model used here draws its basic structure from the DICE-99 model of climate and economy (see Nordhaus and Boyer, 2000). Where the DICE-99 model solves a deterministic problem, our model admits uncertainty over the critical values of the global climate system which, once crossed, lead to a state in which THC shutdown becomes irreversible. Uncertainty is also included over the sensitivity of climate to changes in atmospheric carbon levels.

2.1 The Climate Model

The global climate system is modeled as in Nordhaus and Boyer (2000), with the addition of laws of motion describing the evolution of the THC. The climate module provides a mapping between climatic state today and climatic state in the next period as a function of economic activity (GHG emissions from global production). The schematic diagram in Figure 1 shows an overview of the system. Despite the similarities to Nordhaus and Boyer (2000), the model is presented in its entirety below for clarity of notation.

Climate change is driven by carbon[1] (or equivalently carbon dioxide (CO_2)) mass in the atmosphere.[2] This stock (m) evolves as an autoregressive function of the existing stock (m_{t-1}) and period emissions (E). The law of motion is characterized in (1) with a decay rate ($\delta_m \in (0,1)$) and a preindustrial, baseline carbon mass m_b, as:

$$m_t = E_t + (1-\delta_m)(m_{t-1} - m_b) + m_b. \tag{1}$$

Increasing carbon mass in the atmosphere causes an increase in radiative forcing, or what is commonly known as an enhancement of the greenhouse effect. Global surface (G) and ocean

[1] Note that the DICE-99 model characterizes carbon accumulation as a three-reservoir model. Here, we adopt a characterization used in earlier versions of the DICE model for the reason that the slight loss of accuracy in the transition of carbon stocks allows us to reduce the state space by two variables and thus to significantly reduce computation time.

[2] As long as m_t and E_t are consistent in measuring either carbon or CO_2, these are interchangeable. The emissions:output ratio must also reflect the choice of whether to model carbon or CO_2.

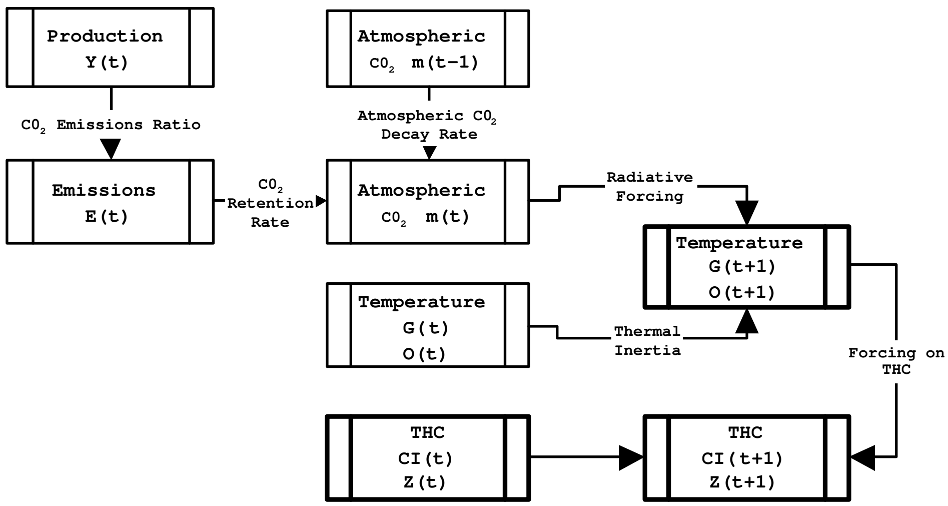

Figure 1: Schematic Diagram of the Climate Model

(O) temperature changes are influenced by the amount of radiative forcing, determined by climate sensitivity parameter (η), and the relative size of the carbon stock, as well as current temperature levels as:

$$G_{t+1} = \lambda_1 G_t + \omega O_t + \eta \frac{\log\left(\frac{m_t}{m_b}\right)}{\log(2)}, \tag{2}$$

and

$$O_{t+1} \quad = \quad \lambda_2 O_t + (1 - \lambda_2) G_t. \tag{3}$$

The parameters λ_1 and λ_2 are the autoregressive parameters in each relationship, while ω represents the buffering effect of the oceans on surface temperature changes. Given this characterization, we can express the long-run temperature change from a doubling of CO_2 as $G_{2\times CO_2} = \frac{\eta}{1-\lambda_1-\omega}$.

In order to model the THC, additional laws of motion are added to the original DICE-99 model. Specifically, we model perturbations in the global climate system as exerting a force on the current circulation, in a way similar to the modeling of the evolution of temperature as a function of CO_2 emissions. The climate pair $(G, \Delta G)$ is assumed to capture the extent of perturbation to the system exerted by climatic changes, where ΔG represents the rate of change of global surface temperature. The circulation rate, ($CI \in [1, 0)$), at any time period is a function of its value in the previous time period and the amount of perturbation in the system. A threshold level of circulation ($\overline{CI}$) is defined such that once circulation falls below this level, a shutdown of the THC becomes irreversible. The absorbing state in which the THC is shutting down is described by the binary state variable $Z \in \{0,1\}$, where 1 indicates the shutdown state. Thus, the THC system is modeled as a two equation system such that:

$$CI_{t+1} = \begin{cases} \nu CI_t + (1-\nu)(1-\xi_1\left(\frac{G_t}{\overline{G}}\right)^{\kappa_1}\left(\frac{\Delta G_t}{\overline{\Delta G}}\right)^{\kappa_2} - \xi_2 G_t^2), \text{ if } Z_t = 0 \\ \xi_3 CI_t, \text{ if } Z_t = 1 \end{cases} \tag{4}$$

$$Z_{t+1} = (1-Z_t)S(CI_{t+1}) + Z_t. \tag{5}$$

The functional form in (4) was chosen to reflect stylized facts from the Stocker and Schmittner (1997) study. Specifically, we needed a characterization which would respond to changes in the temperature level as well as alterations in the rate of change of temperature, however the function must continue to exert negative forcing on circulation when the rate of change of temperature is close to zero. The perturbation term $(G, \Delta G)$ has elements re-scaled by parameters $\overline{G}$ and $\overline{\Delta G}$. The parameters ξ_1 and ξ_2 capture the effect of the total perturbation and the residual effect of a change in temperature level respectively. Parameters κ_1 and κ_2 describe the shape of the response to climate pairs. The parameter $\xi_3 \in (0,1)$ captures the slow, inescapable shutdown of the circulation once the threshold has been crossed. $S(CI_{t+1})$ is a threshold function which returns a value of 0 or 1 depending on whether the climate system has crossed the critical level for a shutdown. We adopt a binary function for $S()$, such that:

$$S(CI_{t+1}) = \begin{cases} 1 \text{ if } CI_{t+1} < \overline{CI} \\ 0 \text{ if } CI_{t+1} \geq \overline{CI}, \end{cases} \tag{6}$$

which is parameterized by the threshold constant $\overline{CI}$.

We assume that global temperature is not affected by the thermohaline circulation. The effect of a shutdown in the circulation would be more accurately characterized by a change in the distribution of thermal energy throughout the globe, with a cooling over the North Atlantic region and increased warming in the Southern Hemisphere; see for instance Vellinga and Wood (2002).

2.2 The Economic Environment

The global economy produces a single good using a Cobb-Douglas technology with inputs of capital and labour. The good is used for consumption and investment. Production in the economy is negatively affected by changes in global climate as well as by the social planner's choice of the level of emissions reduction (μ). Below, we differentiate between the evolution of calendar time (τ) and simulation period t. Calendar time (τ) is expressed in years from a base year, and allows us to capture the evolution of exogenous population and technology trends. The evolution of state variables is determined by the climate model described above, exogenous processes for labour supply and technology, and the planner's decisions on emissions control and investment.

2.2.1 Agents An initial population (measured in billions of people) of L_0 agents identical in all aspects grows through time at a convergent rate, such that population growth eventually limits to zero. The initial growth rate of the population is given by $\gamma_l \in [0,1)$ and the decay of this growth rate toward stability is parameterized by $\delta_l \in [0,1)$ in the law of motion:

$$L(\tau) = L_0 \exp\left(\frac{\gamma_l}{\delta_l}(1 - e^{-\delta_l \tau})\right). \tag{7}$$

Agents in the economy supply labour inelastically for production. Agents in the economy have constant relative risk aversion utility of per capita consumption ($c = \frac{C}{L}$) with parameter σ given by:

$$U(c) = \frac{(\frac{C}{L})^{1-\sigma}}{(1-\sigma)}. \tag{8}$$

2.2.2 Technology Technological change in the model is exogenous, and is described by two variables, A and ϕ, which describe total factor productivity and the emissions efficiency of production. The Hicks-neutral technology parameter A evolves according to:

$$A(\tau) = A_0 \exp\left(\frac{\gamma_a}{\delta_a}(1 - e^{-\delta_a \tau})\right). \tag{9}$$

The emissions output ratio evolves according to the exogenous process given by:

$$\phi(\tau) = \phi_0 \exp\left(\frac{\gamma_\phi}{\delta_\phi}(1 - e^{-\delta_\phi \tau})\right). \tag{10}$$

The parameters γ_a and γ_ϕ denote growth rates in the respective technologies while δ_a and δ_ϕ denote the speed of convergence of these growth rates over time.

It is important to emphasize that exogenous technology levels are a function only of initial conditions and the evolution of calendar time. Thus, with known initial conditions and growth parameters, we can completely describe the state of exogenous technology in the model with knowledge of τ. As such, τ is treated as a state variable in the model, as in Kelly and Kolstad (2001).

2.2.3 Production Production (Y) of the aggregate consumption (C) and investment (I) good occurs according to a Cobb-Douglas production function in two inputs, capital (K) and labour (L) with standard share parameter α given by:

$$Y_t = \frac{(1 - b_1 \mu_t^{b_2})(1 + \theta_3(CI_t - 1))}{(1 + \theta_1 G_t + \theta_2 G_t^2)} A(\tau_t) K_t^\alpha L(\tau_t)^{1-\alpha} = C_t + I_t. \tag{11}$$

Equation (11) includes the cost of chosen emissions control level $\mu \in [0,1)$, the cost of increased surface temperature G, and the cost of decreased THC (CI), and the effect of exogenous technology $A(\tau)$ on the productivity of factors K and $L(\tau)$. The cost of climate change in the economy thus feeds back to agents through consumption. As the climate system, specifically CI and G, deviates from pre-industrial values, or as the planner imposes more stringent emissions control, the productive capacity of the economy decreases.

The parameters b_1 and b_2 are the linear and exponential components of regulatory costs, while θ_1 and θ_2 have the equivalent interpretation for damages from temperature deviations, and θ_3 represents the damages from a complete shutdown in the THC. Abatement decisions are independent across periods, since no abatement capital is required. The right hand side of (11) gives the aggregate budget constraint.

The capital stock (K) evolves according to choice of investment (I) and depreciation (δ_k) through the standard relationship given by:

$$K_{t+1} = (1-\delta_k)K_t + I_t. \tag{12}$$

Emissions in the economy, which represent the input to the climate model, are determined by output (Y), chosen abatement (μ), and emissions efficiency ($\phi(\tau)$) according to:

$$E_t = (1-\mu_t)\phi(\tau_t)Y_t. \tag{13}$$

2.3 Dynamic Optimization

A social planner is assumed to maximize social welfare, given by the discounted sum of utility from per capita consumption by choosing aggregate investment and emissions control subject to a capital accumulation constraint and the laws of motion implied by the climate model. The state vector for the economy is summarized by $\mathbf{S_t} = (K_t, G_t, O_t, m_t, CI_t, Z_t, \tau_t)$, while investment ($I$) and emissions control (μ) are the control variables. Denoting next period variables with a prime, we can express the indirect utility arising from the solution to the social planner's problem by the Bellman equation V, which is characterized as the fixed-point solution to the dynamic program described below:

$$V(\mathbf{S}) = \max_{I,\mu} L(\tau)\frac{(C/L(\tau))^{(1-\sigma)}}{1-\sigma} + \beta E\left[V\left(\mathbf{S}'\right)\right] \tag{14}$$

s.t.

$$C = \frac{(1-b_1\mu^{b_2})(1+\theta_3(CI-1))}{(1+\theta_1 G+\theta_2 G^2)} A(\tau)K^{\alpha}L(\tau)^{1-\alpha} - I, \tag{15}$$

$$K' = (1-\delta_k)K + I, \tag{16}$$

$$G' = \lambda_1 G + \eta\frac{\log\left(\frac{m}{m_b}\right)}{\log(2)} + \omega O, \tag{17}$$

$$O' = \lambda_2 O + (1 - \lambda_2)G, \tag{18}$$

$$m' = (1 - \mu)\phi(\tau)Y + (1 - \delta_m)(m - m_b) + m_b, \tag{19}$$

$$CI' = \begin{cases} \nu CI + (1 - \nu)(1 - \xi_1 \left(\frac{G}{\overline{G}}\right)^{\kappa_1} \left(\frac{\Delta G}{\overline{\Delta G}}\right)^{\kappa_2} - \xi_2 G^2) \text{ if } Z = 0 \\ \xi_3 CI \text{ if } Z = 1, \end{cases} \tag{20}$$

$$Z' = (1 - Z)S(CI') + Z, \text{and} \tag{21}$$

$$\tau' = \tau + 1, \tag{22}$$

where $\beta \in [0, 1)$ is the discount factor applied by the planner to future welfare. The expectations operator in (14) is with respect to the climate sensitivity parameter η and the threshold in $S(CI')$, both treated as unknown to the planner.

3 Calibration

Clearly, the results of the simulations will be highly sensitive to the calibration of the model. In order to derive parameter values which are most consistent with observed historical data and future projections, we adopt a modular approach, beginning with the climate module which we are able to calibrate independent from the economy. The economic model is then calibrated to match emissions and production data and projections for the time period of 1970 to 2030.

In order to calibrate the THC module, we use experimental data from Stocker and Schmittner (1997). Stocker and Schmittner (1997) presents five emissions scenario experiments, and test the response of the thermohaline circulation using a coupled ocean-atmosphere climate model, Bern 2.5-D (Stocker et al., 1992). The five scenarios are stylized emissions control patterns which are defined by the growth rate of the stock of atmospheric CO_2 and a ceiling on its concentration. For example, the scenario labeled as (650 ppm, 1%) indicates that atmospheric concentration grows by 1% until reaching a ceiling of 650 ppm, and remains stable thereafter. The carbon mass profiles (GtC) for each of the scenarios are shown in Figure 2.[3] The predicted temperature evolutions for each of the scenarios (from our climate model) are shown in Figure 3. The simulated paths for ocean circulation developed in Stocker and Schmittner (1997) are shown in Figure 4, while the response functions generated by our climate model are shown in Figure 5.

Parameters for equations (4), (5), and (6) were chosen such that the dynamics of our model climate system closely match the Stocker and Schmittner (1997) simulations. Specifically, we allow a Quasi-Newton (BFGS) routine to choose values for the set of parameters $(\overline{G}, \overline{\Delta G}, \nu, \xi_1, \xi_2, \xi_3, \kappa_1, \kappa_2)$, conditional on $\overline{CI} = .5$ such that simulations of the climate model have the lowest sum of squared differences from the Stocker and Schmittner (1997) paths. The sum of squared residuals, in terms of percentage circulation remaining between

[3] We convert ppm to GtC on the axis of the figure since these units are used throughout the paper.

the modeled pathways and the experimental data, was .3564. The only significant improvement which could be achieved was in decomposing the parameter value for ν in (4), the inertia of circulation, to be ν_- when circulation was decreasing and ν_+ when it was increasing, which allowed for the downward response to be slightly faster than the upward response. Given this change, the scenarios proposed by Stocker and Schmittner (1997) were replicated with an error of .2001, however since this requires the modification of the state variable Z to a tri-modal variable, we decided to keep a single parameter ν in the specification.

The simulations cited above rely on parameters of the climate model which govern temperature changes as a result of changes in atmospheric carbon. The simulations use a climate sensitivity implying 3°C of warming for a doubling of atmospheric carbon. We parameterize (2) by specifying a temperature autoregression parameter (λ_1) of .9122 and a mixing parameter (ω) of .002, which are roughly the annualized equivalents of the Nordhaus and Boyer (2000) parameterization. Note that we also fixed $\overline{CI}$ such that a shutdown becomes irreversible after half of the circulation has deteriorated. This is consistent with the findings of Stocker and Schmittner (1997) and Knutti and Stocker (2002).

The parameter (θ_3) which maps THC decline into productivity is difficult to estimate and difficult to aggregate in a global model. Vellinga and Wood (2002) states that much of the damage from a shutdown in these ocean currents will occur regionally, with cooling in Europe and Northeastern North America, partially compensated by a warming of the Southern Hemisphere. Tol (1998) states that these costs could amount to between a 0% and a 3% productivity loss in Western Europe, and as such, Keller et al. (2004) adopts an expected value of 1.5% for this parameter. Mastrandrea and Schneider (2001) conducts a sensitivity analysis with damages ranging from 1% to 10%. In light of the wider range considered by Mastrandrea and Schneider (2001), we choose a parameter value of 2.5% in this analysis.

Finally, we must calibrate the levels of uncertainty to be faced by the social planner over the threshold level and the climate sensitivity to carbon levels. We assume that the planner views these parameters as normally distributed, independent, random variables. In order to fix a confidence interval around the parameter governing the sensitivity of climate to CO_2 levels, we use the range proposed by the IPCC (2001) of 1.5°C to 4.5°C for a doubling of CO_2 relative to preindustrial levels, which is consistent with the scientific confidence interval discussed in Knutti et al. (2002). For the analysis presented below, uncertainty over climate sensitivity is achieved by setting a value of the parameter standard deviation (σ_η) that implies a stable 95% confidence band admitting 3 ± 1.5°C for G_{2xCO_2}.[4] For the threshold level below which shutdown becomes irreversible, we choose a mean value of .5, and set a 95% confidence interval of $.5 \pm .1$ for the critical fraction of remaining circulation

[4] The normal distribution allows for the planner to admit the possibility of a much higher climate sensitivity, but with a correspondingly low probability.

through setting the parameter standard error $\sigma_{\overline{CI}}$. This choice is consistent with the findings of Knutti and Stocker (2002) that the THC does not recover after crossing a 50% reduction, and Manabe and Stouffer (1993), which shows a recovery of the THC after a reduction to less than 40% of its full value.

The first step in calibrating the economic model is to set the exogenous trends for labour supply and the emissions:output ratio. We calibrate the evolution of the emissions:output ratio by choosing the parameter values for equation (10) which match observed and predicted values for global economy from the International Energy Agency (IEA) (2004). The fit of the exogenous trend to this data is shown in Figure 6. The labour supply is calibrated to the median scenario from United Nations (2004) data and projections on the global population, and the fit is shown in Figure 7.

We then fix most parameters of the economic sector to values used in Pizer (1999) and Nordhaus and Boyer (2000). In order to match the growth characteristics observed in the IEA (2004) data, we choose a rate of exogenous technology growth such that two conditions are met. First, we want the economy under a business-as-usual assumption of no climate change policy from 1970-2030 to match IEA (2004) data on global GDP. Second, in order for the solution to the dynamic program described in (14) to exist, all exogenous technology variables must be bounded within the confines of a grid, and as such we must calibrate the law of motion for A such that it converges to a constant. We choose a value of δ_a such that technological change ceases after 800 years, and then choose γ_a to match observed economic growth conditional on this assumption. The fit of the model to global GDP is shown in Figure 8. Combined with the calibration of the emissions:ouput ratio (ϕ) above, we are thus able to replicate the dynamics of global emissions as shown in Figure 9.

The values for each of the parameters in the model are shown in Table 1, while economy and climate 1970 starting values are in Table 2. Readers interested in a complete discussion of the calibration of economic and climate sectors of integrated assessment models with uncertainty are directed to Nordhaus and Boyer (2000), Pizer (1999), or Leach (2004).

4 Solution

The solution to the social planner's problem proposed in equations (14)-(22) is described below. The value function for the social planner's problem is solved using a fixed point algorithm which makes use of a multi-layer, feed-forward neural network solved as a non-linear regression function to approximate the value function, as outlined in Kelly and Kolstad (2001).[5]

Briefly, the solution approach involves using a neural network as a flexible functional form to approximate the value function in equation (14). Recall that **S** represents the

[5] Solution and simulation calculations performed using Ox version 3.30, Doornik (2003).

Table 1: Calibrated Values under Uncertainty

Parameter	Description	Calibrated Value
Economic Parameters		
β	Discount factor	.9613
δ	Capital depreciation rate	.045
σ	Coefficient of relative risk aversion	1.3
α	Production share of capital	.3
A_0	Initial exogenous technology	0.01612
γ_a	Initial growth rate of labour augmenting tech. change	0.00955
δ_a	Decay rate on γ_a	.004
L_0	Initial population	3100
γ_n	Initial growth rate of population	0.024
δ_n	Decay rate on γ_n	.017
ϕ_0	Initial emissions:output ratio	0.303402
γ_ϕ	Initial growth rate of emissions:output ratio	-0.02425
δ_ϕ	Decay rate of emissions:output ratio	.02042
b_1	Linear emissions control costs	.0690
b_2	Exponential emissions control costs	2.877
Carbon Cycle and Climate Parameters		
m_b	Preindustrial concentration of CO_2 (GtC)	590
$(1-\delta_m)$	atmospheric retention of CO_2	.99167
λ_1	Autoregressive component of temp. change	.9112
λ_2	AR(1) parameter in ocean temperature	.9912
$G_{2\times CO_2}$	Temperature sensitivity to CO_2 doubling (°C)	3
ω	Coefficient on ocean temp. in temperature Change	.0002
THC-specific Parameters		
$\overline{G}$	Threshold Level of G for THC shutdown(°C)	7.81354
$\overline{\Delta G}$	Threshold level of ΔG for THC shutdown (rate of change)	1.2
$\overline{CI}$	THC Reduction at threshold	0.5
ξ_1	Scaling factor in THC reduction	381.314
ξ_2	Linear shift for steady state THC reduction	.01115
ξ_3	Implied half-life in years of THC in shutdown state	50
κ_1	Exponent on temperature in THC reduction	1.3298
κ_2	Exponent on rate of temperature change in THC reduction	1.3584
ν	AR1 component in THC reduction	.977989
Climate Change Damage Parameters		
θ_1	Linear component of damages	.00071
θ_2	Exponent in damage function	.00242
θ_3	Loss in GDP from shutdown in THC	.025
Uncertainty Parameters		
σ_η	Standard deviation of η	0.0860472
$\sigma_{\overline{CI}}$	Standard deviation in $\overline{CI}$ for THC shutdown	.05102

Table 2: State Variables Used in the Model

State	Definition	Units	Grid Min	Grid Max	Starting Value (BM)
K	Capital stock	10^{12} \$1987	1	2500	5
m	Atmospheric CO_2	10^9 tonnes	590	3500	678
G	Global surface temperature	°C from mean	0.2	6	0
O	Ocean temperature	°C from mean	0.2	6	0
CI	Thermohaline circulation	% remaining	0	1	1
Z	Indicator for absorbing shutdown state (Z=1)		0	1	0
τ	Calendar time (state of technology and population)	years	1	1000	1

state vector, and the control variables are emissions control (μ) and investment (I). Let ϱ represent the set of model parameters and define $h(\hat{\varrho})$ as the probability density function for a given set of model parameters, conditional on the planner's uncertainty over $\overline{CI}$ and η.[6] Denote by $g(S, \mu, I, \varrho)$ the transition function, capturing the evolution of the state variables as a function of their current values, controls, and parameters of the model in equations (16-22). Where V_i represents iteration i of the value function, we know from the Bellman principle that we can express the indirect utility across states under the solution to the planner's problem as the fixed point of:

$$V_i(\mathbf{S}) = \max_{\mu,I} U(C|S, \mu, I) + \beta \int_{\hat{\varrho}} V_{i-1}(g(S, \mu, I, \hat{\varrho})h(\hat{\varrho})d\hat{\varrho}. \tag{23}$$

The use of a neural network provides a flexible, non-linear functional form which can be trained to match very closely the unobservable value function of the social planner. Specifically, we use the following functional form to map the value function across N grid points in the k dimensional state space:

$$\Phi(\mathbf{S_j}) = \chi_1' \tanh\left(\chi_2 \mathbf{S_i} + \chi_3\right) + \chi_4. \tag{24}$$

Here, $\mathbf{S_j}$ is the $k \times 1$ vector of state variables, χ_1 is an $l \times 1$ vector of weights, χ_2 is an $k \times l$ matrix of weights, χ_3 is an $l \times 1$ vector, and χ_4 is a constant. l is described as the number of nodes in the network, and determines the fineness of the approximation. As in Kelly and Kolstad (2001), l is chosen such that it is the closest integer less than $\frac{(N^{0.5}-1)}{(k+2)}$.

We deviate from the solution algorithm employed by Kelly and Kolstad (2001) only slightly, in that we estimate the value function over two of the exogenous technology

[6] For parameters over which there is no uncertainty, $h(\hat{\varrho}) = 0$ where $\hat{\varrho}$ includes any value for that parameter other than the true value.

variables as well as the time period. Since A and ϕ are each exponential functions of the time period, they can be included as separate elements of S for fitting the neural network function, and they are important directly in the return to both control variables. The second deviation is that, as opposed to choosing discrete values for each variable and creating a grid as the cross-product of these discrete values, we draw states from a k-dimensional hypercube using a low-discrepancy sequence (see Judd, 1998).

The solution algorithm which solves the fixed point problem (23) using $\Phi(S)$ as an approximation to the value function is provided below:

Algorithm 1
Objective: Compute an approximation to the value function using a recursive, iterative algorithm and approximation by neural network.

Algorithm Preliminaries: Define ranges[7] *for each of the state variables to be covered by a grid, and a number of points N to make up the grid. Draw N points from a k-dimensional, low discrepancy sequence. Choose a convergence criterion ϵ and an initial guess of the value function. We use the welfare from consumption of production in the first period as a guess to train the neural network.*

Initialization: Estimate parameters χ^0 of equation (24) as an approximation to the value function.

Step 1: Using (14), calculate the optimal choice variables and the expected value at each point in the state space. Use analytic derivatives of the approximated value function (Φ) to generate a system of two first order conditions as zero functions in the two unknowns I and μ. Solve for I and μ taking account of potential corner solutions.

Step 2: Letting i represent the number of iterations of Step 1 and Step 2 thus far completed, estimate new parameters χ^i of equation (24) as the updated approximation to the value function, $\Phi^i(S)$ using values generated in Step 1.

Step 3: Repeat Step 1 and 2 until convergence criteria ϵ satisfied for the value function.

The algorithm used in this paper relies on the parameters of $\Phi(\mathbf{S})$ as the global approximation of the period value function. In order to evaluate the integrals in the planner's problem, nested Monte Carlo methods are used. Specifically, we use 5000 grid points, and set the number of nodes in the neural network such that $l = 5$. We are able to fit the value function with a sum of squared residuals on the order of $5 * 10^{-5}$. Within each calculation, we use 150 iterations to calculate the Monte Carlo integrals for the expected value over future state transitions.

[7] Maximum and minimum values for each state variable on the grid are shown in Table 2.

Three versions of the model are solved using the above algorithm, and will be discussed further in the simulations below. First, we solve a model with no uncertainty, setting σ_η, and $\sigma_{\overline{CI}}$ each to zero. Second, a model with uncertainty only over the critical circulation threshold, $\overline{CI}$, setting $\sigma_{\overline{CI}}$ to the value shown in Table 1 is solved. Finally, the full model with uncertainty added by setting σ_η and $\sigma_{\overline{CI}}$ to the values shown in Table 1 is solved. We will refer to these model parameterizations below as UC0, UC1, and UC2 respectively. It is most important to note that the transitions over state variables are determined in simulation.

5 Simulation

The results of this analysis are captured through simulation of the optimal policy model as well as characterization of the estimated value function. Once we have solved the three versions of the model described above, we are able to define the indirect utility of any point in the state space, and as such we are able to determine the cost of uncertainty and the value to agents of the imposition of emissions abatement policy in the presence of uncertain thresholds. For each of the simulations, we simulate the economy using the same starting values as in the calibration section, however we constrain the economy to follow the no-abatement benchmark until 2005.

In Subsection 5.1, we undertake a policy analysis approach, in addition to reporting optimal policy results. Below, we first examine the characteristics of the optimal policy in the sense implied by the UC0 model. We compare this policy choice first to an imposed no-abatement benchmark, and second to the imposition of policies with the goal of preventing a shutdown of the THC.

In Subsection 5.2, we examine how model predictions and welfare are altered by the addition of uncertainty in the decision making framework in UC1 and UC2. We examine the indirect utility costs of uncertainty, and attempt to add context to these results by analyzing the estimated value function.

5.1 Cost-benefit analysis

5.1.1 The optimal policy Optimal emissions control policy in the model consists of a sequence of control rates which is shown for model UC0 in Figure 10. We use the same economy that was simulated in the calibration section, however we constrain emissions control to zero until the year 2006. The imposition of these control rates reduces the quantity of emissions from production, and as such mitigates global climate change. The simulated evolution of the THC and surface temperature under the optimal policy sequence and under a no-abatement benchmark are shown in Figures 12 and 13.

The model predicts that with full information and an optimal emissions policy, the THC collapse threshold will be crossed in 2107, as opposed to 2086 under the no-abatement benchmark. The fact that the optimal policy allows a shutdown to occur is likely a consequence of the distance in time after which a full shutdown occurs. In the last period of the simulation, we still see approximately 25% of the circulation remaining, such that the shutdown produces approximately a maximum 2% damage rate 200 years in the future.

To contrast this result with results in the literature, Mastrandrea and Schneider (2001) predicts that, for 10% collapse-specific damages, a planner will be prepared to allow a collapse if the discount factor is greater than 1.5%. Keller et al. (2000) finds that if threshold-specific damages are .86% or greater, THC preservation policies would be justified (with a discount rate of 3%). We show this result as consistent with the optimal policy for a discount factor of 4% and damages of 2.5%.

5.1.2 The cost of inaction In the optimal regulation model, a social planner chooses both consumption and abatement in each time period, while in the benchmark model, only consumption is chosen, and abatement is constrained to zero in all periods. In order to measure the cost of inaction, we report the shadow value of the constraint as follows. In each period, suppose that a higher authority forces the social planner to delay by one period the imposition of the optimal policy, and then asks consumers to quantify how much of their consumption in that period they would be willing to give up in return for having the optimal policy imposed today as opposed to one year from now. We can measure the shadow value of the restriction by comparing the indirect utility of climate and economic states along each transition path. In Figure 11 we show the compensating variation for the difference in indirect utility obtained on the transition path under the optimal policy and constrained benchmark. As expected, the constraint becomes more costly through time, and we can see that agents are eventually willing to give up as much as 30% of their current consumption in order to impose the optimal policy.

5.1.3 The cost of overreaction: THC preservation policies The results of the optimal policy simulation show that the policy maker will clearly find it sub-optimal to reduce emissions sufficiently to avoid the damage brought on by a shutdown in the THC. Since we are interested in quantifying the welfare losses of a policy preventing a THC collapse, we propose two simulations where additional regulation is undertaken to keep the global carbon levels below 560 ppm and 650 ppm respectively. These numbers are chosen since the Stocker and Schmittner (1997) experiments predicted a stable THC under similar concentration caps. The model replicates these predictions, with circulation rebounding slowly after the corresponding emissions constraints are put into place, as shown in Figure 15. The economic costs of these policies are large, as shown in Figure 14, relative to the optimal policy choice. Welfare effects of this policy change are very small. The discounted utility of per capita consumption for a representative agent from the year 2000 forward will be .01% lower under the 560 ppm scenario and .04% lower under the 650 ppm scenario than under the optimal policy. What we are seeing here is the product of three assumptions in the model. First, since the concentration target is reached far in the future, the welfare effects are small due to any differences in consumption being discounted at over 4%. Second, the costs of control are stipulated as in Nordhaus and Boyer (2000), such that the most costly period of emissions control has less than a 3% productivity cost. Finally, the difference in THC damage in the two economies is, at most 1.4% in the 650 ppm scenario and 1.6% in the 560 ppm scenario.

We can also ask a similar question about a "Kyoto Forever" scenario which restricts global emissions to a 6% reduction over 1990 levels, from the year 2010 onwards. This differs from the above simulations since it forms a constraint on emissions, not concentrations, and represents a much more aggressive policy with earlier onset. In terms of effect on the THC, clearly the restriction has the effect of insuring a stable circulation, as shown in Figure 17, but with significant economic costs, shown in Figure 16, relative to the UC0 optimal policy benchmark. The welfare (discounted utility of consumption) for a representative agent from the year 2000 forward will be 0.14% lower under the Kyoto scenario than under the optimal policy.

Clearly, each of these sub-optimal policies seek to reduce the effects of climate change without taking account of engendered costs and benefits. We therefore expect there to be welfare costs associated with these policy choices relative to the optimal. It is somewhat counterintuitive that they should be as small as they are, although this can be clearly attributed to parameters of the model.

5.2 The costs of uncertainty

A key goal of the paper was to address the cost of uncertainty. Since we have an estimate of the value function in the benchmark case as a metric for the full-information optimum, we can evaluate the indirect utility (the discounted utility from all future consumption conditional on optimal decision making) for any combination of state variables with zero uncertainty. As such, we have a natural measure of the cost of uncertainty, or the value of full information. Along the transition path in the uncertainty case (UC1 or UC2), we can evaluate the difference between the uncertainty value function and the full information value function at each state. This represents the difference between the expected present value of future utility under uncertainty and the same measure under certainty, beginning from the same economic state. The welfare costs to uncertainty are small. We find that there is always a positive value of information, however the welfare differences are quite small. For the UC1 scenario, the value of information along the transition path has a maximum value of a .0005% increase in the welfare measure used in the social planner's problem. For UC2, the value of full information is slightly larger, with a maximum value of .011%. The value of information across the UC1 and UC2 transition paths is shown in Figure 18. The effect of uncertainty here is clearly quite small, contrary to conclusions reached in Keller et al. (2004) that the effect of uncertainty is large enough to induce changes in abatement policy.

Intuitively, we can expect uncertainty to have two countervailing effects. First, uncertainty over climate sensitivity and potential thresholds distorts the planner's decision in whether to invest in emissions control or future capital. A high value of climate sensitivity or an easily-reached threshold imply higher returns to emissions control, and lower returns to future capital, all else being equal. Given greater uncertainty over the returns to emissions control than over the return to future capital, we would expect to see under-regulation. This motive is counter-balanced by the precautionary motive implicit in the utility function, which is likely to lead to over-regulation. We see in Figure 19 that the differences in emissions control policy chosen here are negligible in the face of uncertainty.

In order to provide intuition to these results, we present some graphical representations of the value function. Specifically, we take simulated state variable values at 2005 and 2100 as fixed for all but the variable of interest, which we allow to vary. The corresponding 2-dimensional representation gives us a sense of the curvature of the future value function with respect to evolution of the variable of interest, and as such the costs to uncertainty in that evolution.[8] It is worth noting that, although consumption per capita is increasing, the welfare measure of the sum of utility of per capita consumption will be lower in 2100 than in 2000.

In our model, damages are directly proportional to circulation, and not directly affected by crossing the threshold. This produces an indirect utility function which is decreasing but almost linear in circulation, as demonstrated in Figure 20. In this case, uncertainty over the value of circulation next period is not going to be very costly for the planner, since the expected value under uncertainty is very close to the value at the expected circulation.

The second margin over which uncertainty is introduced is over the evolution of future temperatures in response to accumulated carbon in the atmosphere. Again here, we find that the value function is nearly linear in temperatures, as shown in Figure 21.

Where the costs of uncertainty will have an effect is in the evolution of future capital. Figure 22 shows the value function in capital space. Clearly, the curvature of this function will drive most of the costs of uncertainty, since the return to future capital is affected by the evolution of the climate system. However, this is a second order effect, since the value of future capital is known to (and chosen subject to constraints by) the social planner. The production possibility frontier of the economy is affected by the climate, but at a long time horizon, so the effects of climatic change uncertainty on expected future capital stocks are small, and thus this does not provide a propagation mechanism for the cost of climate uncertainty.

[8] Costs to uncertainty arise as a result of curvature in the value function, which increases the difference between the expected indirect utility of next period's state and the indirect utility of next period's expected state. The greater the curvature of the value function, the greater the divergence between these two, and thus the greater the cost of uncertainty.

We can also observe a second counterintuitive result where the social planner may, over part of the transition path, have perceived welfare levels which seem to suggest a benefit to uncertainty, ie. the expected value of future utility under UC1 may be higher than the known value of future utility under UC0. This can be explained by the fact that where the planner is uncertain about having crossed the threshold, she may still consider that the economy can recover and thus has a higher expected future value. In the full information case, she would weigh only the future utility under a shutdown, which will produce a lower value. Once the economy has crossed the threshold, its transition has already been fully determined, but this is not known to the planner without full information. We do see a type of "ignorance is bliss" result in the time periods where the threshold has been crossed but the planner does not know this for sure.

6 Conclusions

The first goal of this paper was to attempt to extend the treatment of extreme events in the economic modeling of climate change to include additional stylized facts from the scientific literature. First, it is generally accepted that there exists a threshold such that a perturbation of the climate sector that is beyond this boundary may lead to an irreversible state. We have modeled a threshold which depends not just on cumulative climate change, but also the rate of change in the state of climate. A further element which has been added to this model is the notion that, once the threshold is crossed, the shutdown does not occur immediately, but rather slowly over a number of decades. These elements add important richness to the model. We have calibrated our model to match as closely as possible a series of THC scenarios generated through a coupled atmosphere-climate model.

We find that, contrary to the conclusions or assumptions of many papers in the literature, a policy to prevent a THC shutdown may be suboptimal for the planner. We are able to quantify the costs to such a policy, finding that this would impose a slight reduction in welfare as defined by the social planner in the full information case.

As noted by the IPCC, "future unexpected, large and rapid climate system changes (as have occurred in the past) are, by their nature, difficult to predict" (IPCC, 2001). We specifically attack the nature of the problem by modeling and quantifying the cost of uncertainty. We find that there may not be significant economic costs to uncertainty in this case as a result of discounting and that these costs are further mitigated if we allow a smooth transition to the absorbing state, and postulate that damages are a function of remaining circulation, not a direct consequence of crossing the threshold.

In the economics literature, we often reduce the treatment of thresholds in climate to binary effects. While thresholds may have such properties in nature, it is clear that the mapping of these thresholds into economic damages is not likely to be as abrupt. Thus, we may be greatly overstating the costs of uncertainty by not examining the nature of damages directly, and implicitly assuming a role in the utility function for a stable climate. Further work in this area should thus be directed at establishing where, if at all, thresholds exists in damages as well as in climate systems.

References

[1] Cubasch, U., and G. Meehl (2001): *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*chap. Projections of Future Climate Change. Cambridge University Press, Cambridge.

[2] Doornik, J. A. (2003): *Ox Version 3.30*. Oxford.

[3] Intergovernmental Panel on Climate Change (2001): *Climate Change 2001*. Cambridge University Press, Cambridge, UK.

[4] International Energy Agency (2004): *World Energy Outlook: 2004*. IEA Publications, Paris.

[5] Judd, K. L. (1998): *Numerical Methods in Economics*. Massachusetts Institute of Technology, Cambridge, Mass., USA.

[6] Keller, K., B. M. Bolker, and D. F. Bradford (2004): "Uncertain climate thresholds and optimal economic growth," *Journal of Environmental Economics and Management*, 48(1), 723–741.

[7] Keller, K., K. Tan, F. Morel, and D. Bradford (2000): "Preserving the ocean circulation: Implications for climate policy," *Climatic Change*, 47, 17–43.

[8] Kelly, D. L., and C. D. Kolstad (2001): "Solving growth models with an environmental sector," *Journal of Computational Economics*, 18, 217–235.

[9] Knutti, R., and T. Stocker (2002): "Limited predictability of the future thermohaline circulation close to an instability threshold," *Journal of Climate*, 15, 179–186.

[10] 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.

[11] Leach, A. (2004): "The Welfare Implications of Climate Change Policy," Submitted Manuscript.

[12] Manabe, S., and R. Stouffer (1993): "Century-scale effects of increased atmospheric CO_2 on the ocean-atmosphere system," *Nature*, 364, 215–218.

[13] 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.

[14] Mastrandrea, M., and S. Schneider (2001): "Integrated assessment of abrupt climatic changes," *Climate Policy*, 1, 443–449.

[15] National Research Council (U.S.), Committee on Abrupt Climate Change (2002): *Abrupt Climate Change: Inevitable Surprises*. National Academy Press, Washington, DC.

[16] Nordhaus, W. D. (1994): *Managing the Global Commons: The Economics of Climate Change*. MIT Press, Cambridge, Mass.

[17] Nordhaus, W. D., and J. Boyer (2000): *Warming the World*. MIT Press, Cambridge, Mass.

[18] Pizer, W. A. (1999): "The optimal choice of climate change policy in the presence of uncertainty," *Resource and Energy Economics*, 21, 255–87.

[19] Schmittner, A., and T. Stocker (1999): "The stability of the thermohaline circulation in global warming experiments," *Journal of Climate*, 12, 1117–1133.

[20] Smith, J., H. Schellnhuber, and M. Mirza (2001): *Climate Change 2001: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Third Assessment Report of the Intergovernmental Panel on Climate Change*chap. Vulnerability to Climate Change and Reasons for Concern: A Synthesis. Cambridge University Press, Cambridge.

[21] Stocker, T., and A. Schmittner (1997): "Influence of CO_2 emission rates on the stability of the thermohaline circulation," *Nature*, 388, 862–865.

[22] Stocker, T., D. Wright, and L. Mysak (1992): "A zonally averaged, coupled ocean-atmosphere model for paleoclimate studies," *Journal of Climate*, 5, 773–797.

[23] Tol, R. S. (1998): "Potential slowdown of the thermohaline circulation and climate policy," Discussion Paper DS98/06, Institute for Environmental Studies Vrije Universiteit Amsterdam.

[24] United Nations Population Division (2004): *World Population Prospects Population Database*. United Nations, New York, USA.

[25] Vellinga, M., and R. A. Wood (2002): "Global climatic impacts of a collapse of the Atlantic thermohaline circulation," *Climatic Change*, 54, 251–267.

[26] Wunsch, C. (2002): "What is the thermohaline circulation?," *Science*, 298, 1179–1180.

[27] Zickfeld, K., and T. Bruckner (2003): "Reducing the risk of abrupt climate change: Emissions corridors preserving the Atlantic thermohaline circulation," *Integrated Assessment*, 4, 106–115.

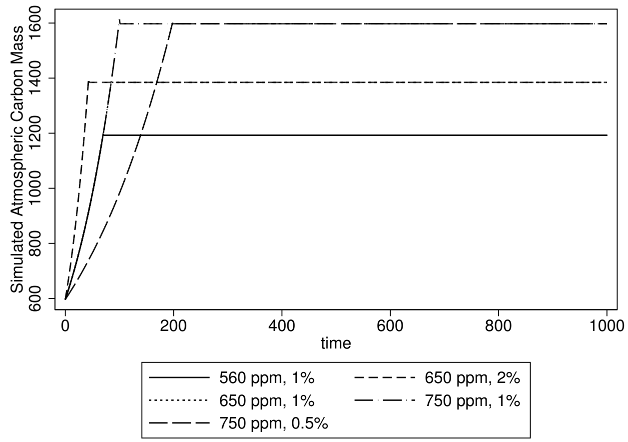

Figure 2: Emissions scenarios, Stocker and Schmittner (1997). Scenarios show evolution of carbon mass in GtC.

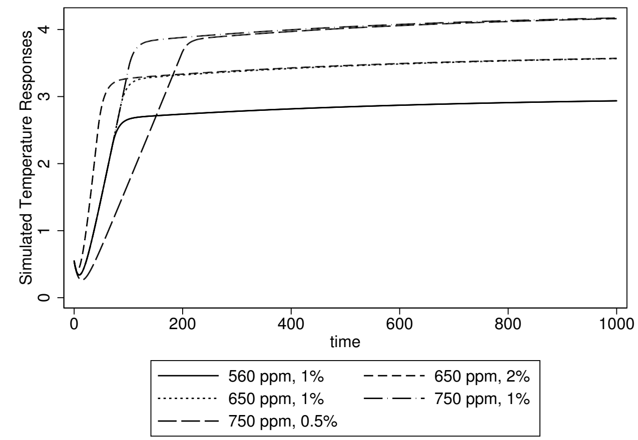

Figure 3: Temperature responses to emissions scenarios, given parameters of the climate model and $G_0 = .55°C$.

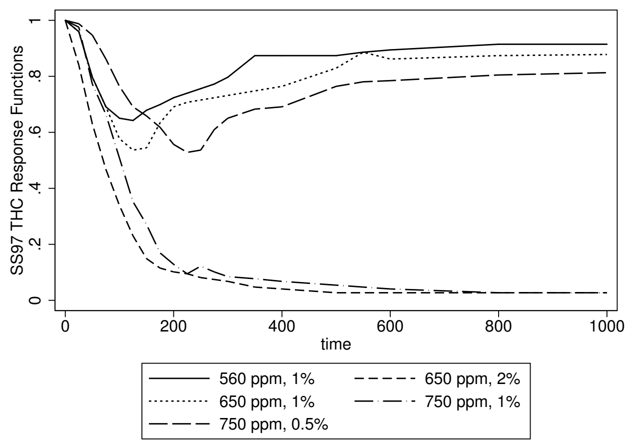

Figure 4: Thermohaline circulation response functions, Stocker and Schmittner (1997). Graph shows percentage of circulation remaining relative to observed first period value.

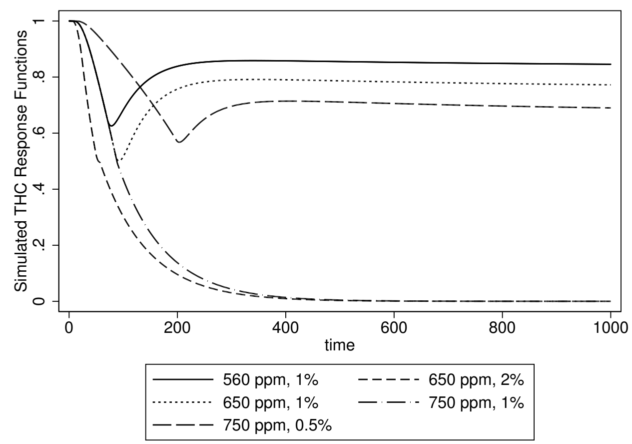

Figure 5: Simulated thermohaline circulation response functions, given parameters of the climate model and temperature evolution.

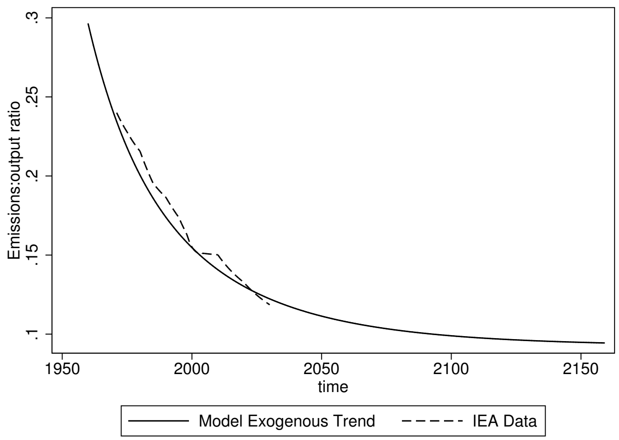

Figure 6: Exogenous emissions:ouput ratio and International Energy Agency data and projections over time.

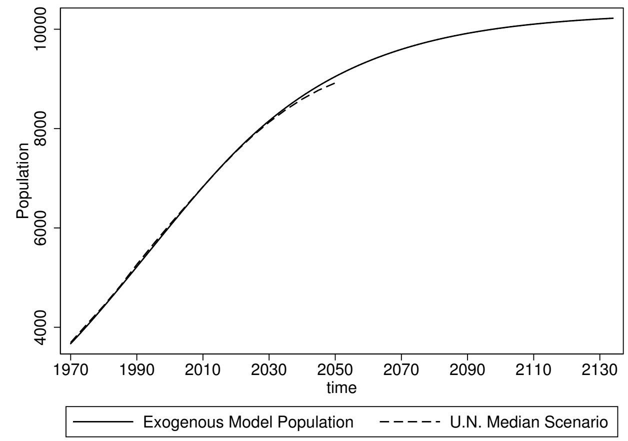

Figure 7: Model and UN projected population (millions of people).

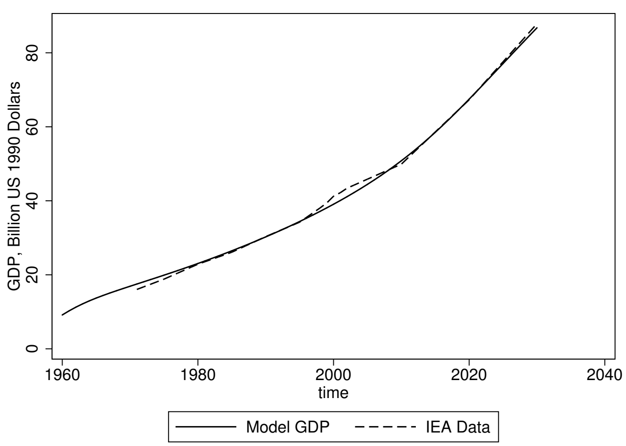

Figure 8: Simulated global GDP versus IEA predictions, UC0 model, no abatement benchmark.

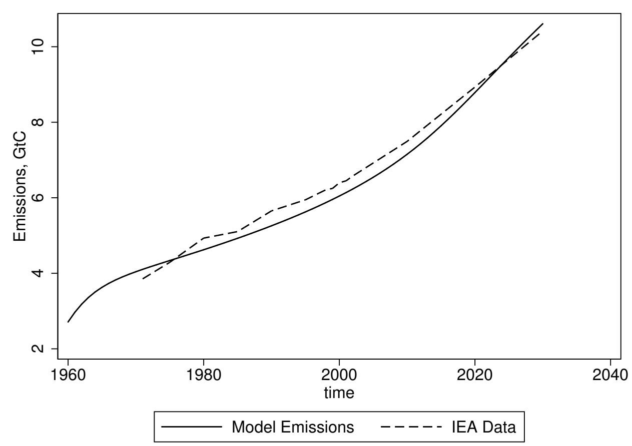

Figure 9: Simulated global CO_2 emissions versus IEA predictions, UC0 model, no abatement benchmark.

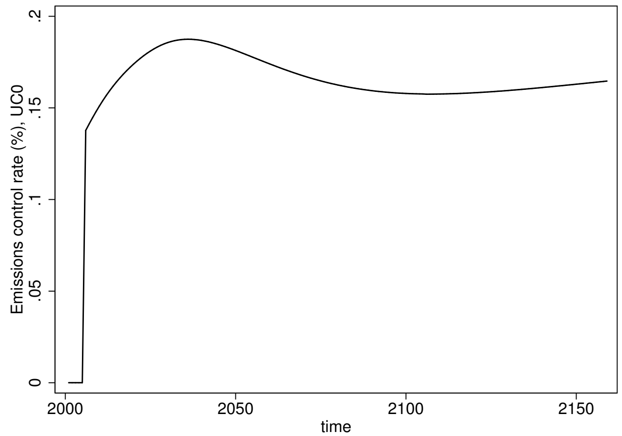

Figure 10: Marginal emissions control rate, optimal policy sequence in model UC0.

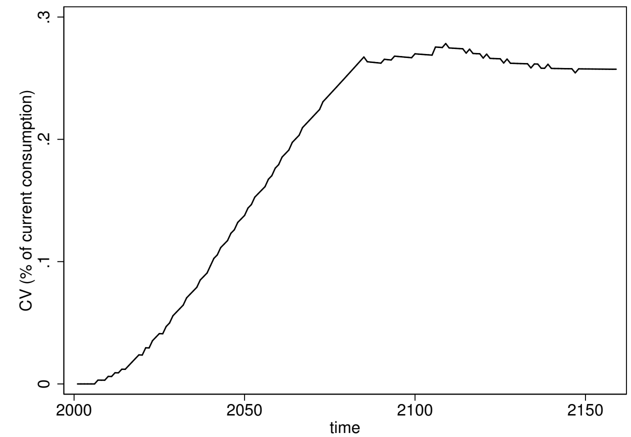

Figure 11: Compensating variation, constrained emissions control, UC0 model. Shows the welfare cost of forcing the no abatement benchmark (BM).

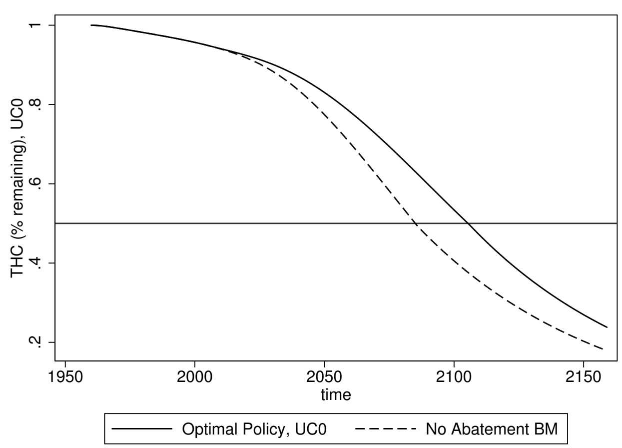

Figure 12: Circulation, optimal and constrained emissions control.

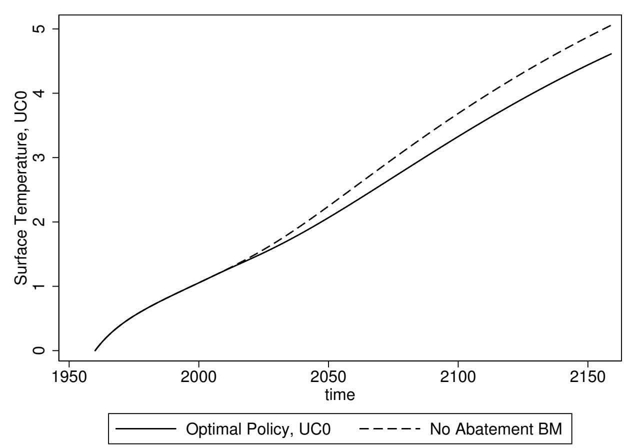

Figure 13: Surface temperature, optimal and constrained emissions control.

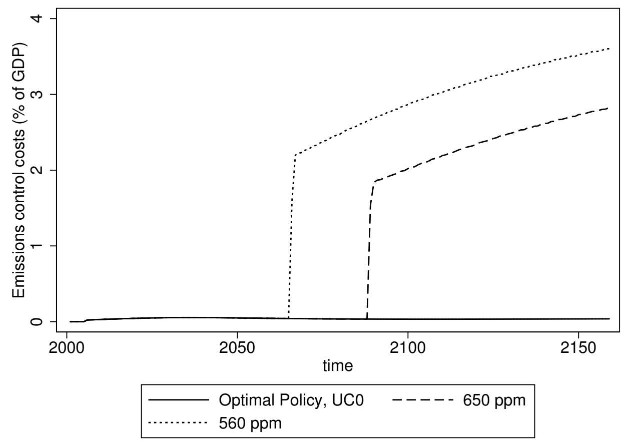

Figure 14: Control costs, 560 ppm and 650 ppm carbon constraints, UC0 model.

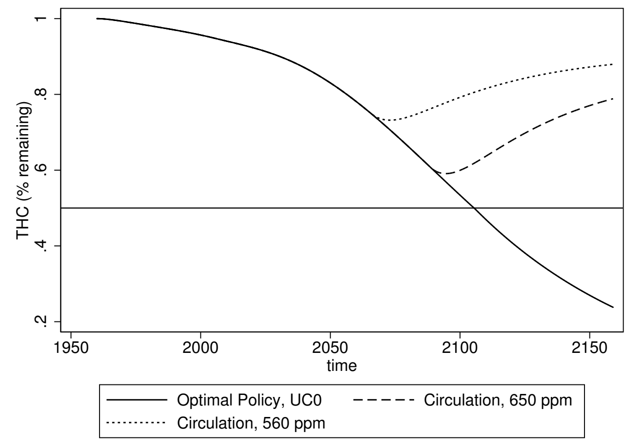

Figure 15: Circulation, 560 ppm and 650 ppm carbon constraints, UC0 model.

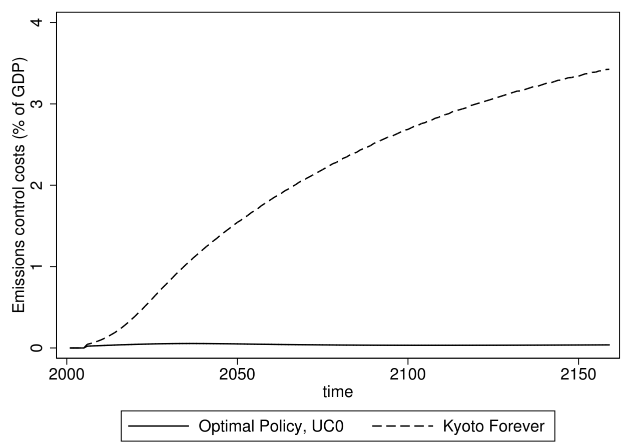

Figure 16: Regulatory cost, Kyoto emissions constraint, UC0 model.

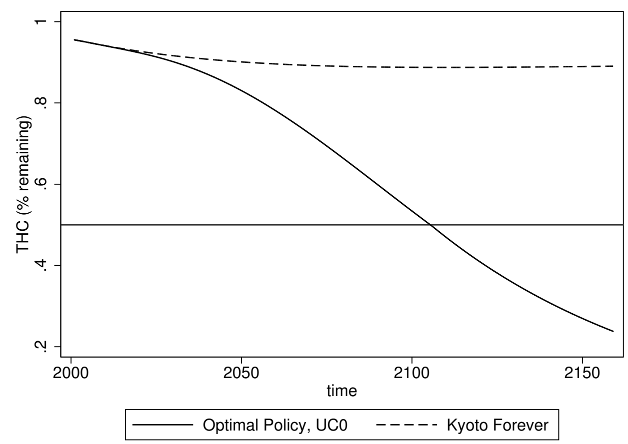

Figure 17: Circulation, Kyoto emissions constraint, UC0 model.

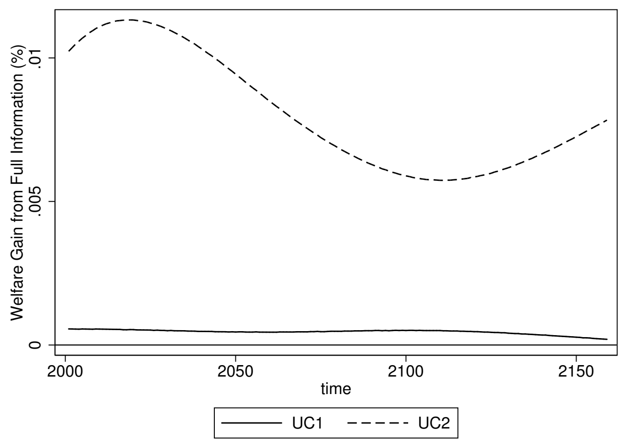

Figure 18: Value of full information under uncertainty. Shows differences in indirect utility of the social planner (welfare) in (%) under full information versus uncertainty over the transition path for UC1 and UC2

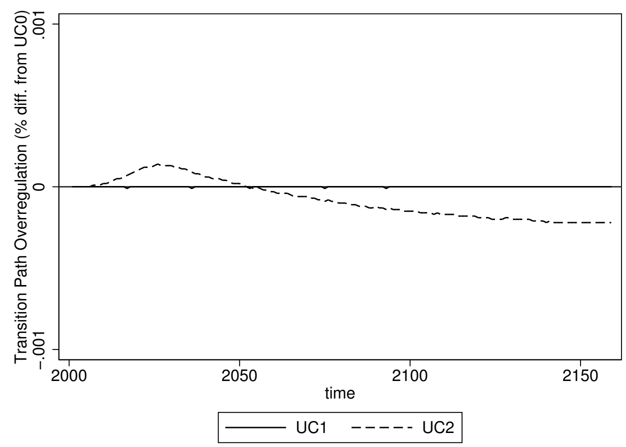

Figure 19: Uncertainty-induced overregulation. Difference between emissions control rate under full information and under UC1 and UC2 respectively.

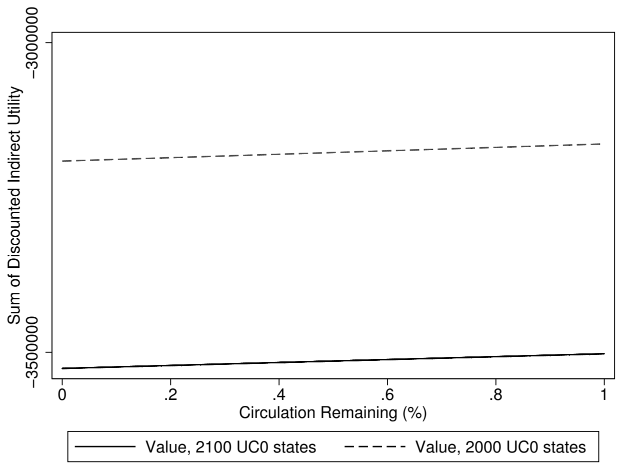

Figure 20: Cross section of welfare over circulation, year 2000 and 2100 under UC0

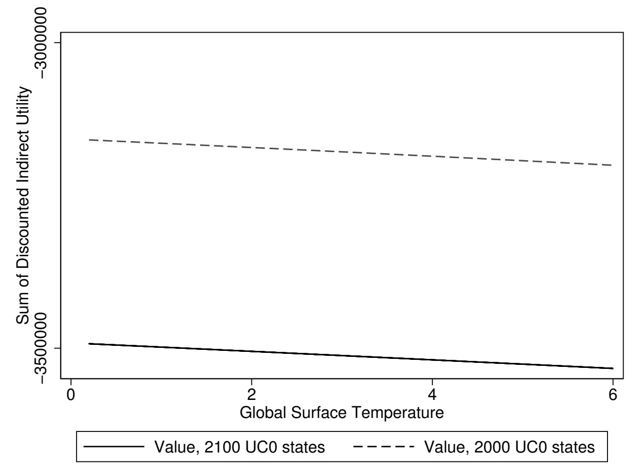

Figure 21: Cross section of welfare over surface temperature, year 2000 and 2100 under UC0

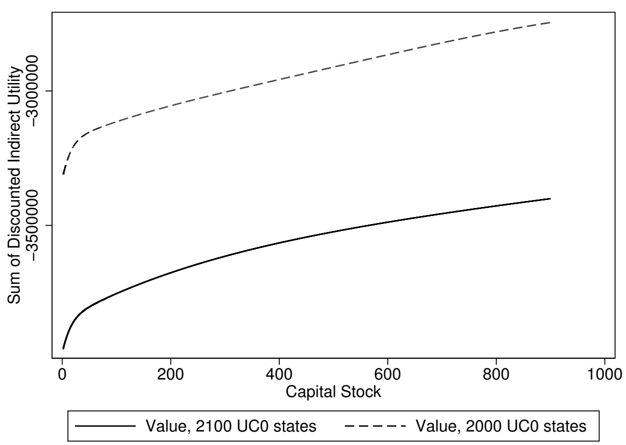

Figure 22: Cross section of welfare over capital stock, year 2000 and 2100 under UC0