Optimal Management of a Stochastic Stock Pollutant with Non-convex Feedback: An Application to Climate Change*

Charles Séguin†

Job Market Paper
01/20/12

Abstract

Non-convexities in feedback processes are increasingly found to be important in the climate system. To evaluate their impact on the optimal greenhouse gas (GHG) abatement policy, I introduce non-convex feedbacks in a stochastic pollution control model. I numerically calibrate the model to represent the mitigation of greenhouse gas (GHG) emissions contributing to global climate change. This approach makes two contributions to the literature. First, it develops a framework to tackle stochastic non-convex pollution management problems. Second, it applies this framework to the problem of climate change. This approach is in contrast to most of the economic literature on climate change that focuses either on linear feedbacks or environmental thresholds. I find that non-convex feedbacks lead to a decision threshold in the optimal mitigation policy, and I characterize how this threshold depends on feedback parameters and stochasticity.

*I thank Larry Karp and Christian Traeger for their invaluable guidance as advisers throughout this project. All errors and omissions are of course mine.

†Ph.D. candidate, department of Agricultural and Resource Economics, University of California, Berkeley. Email: scharles@berkeley.edu

1 Introduction

Optimal control of greenhouse gas (GHG) emissions is an ideal problem to apply the economic theory of stock pollutant management. The benefits from emissions, arising from consumption and production, must be traded off with the damages from changes in the climate, arising from the accumulation of GHG in the atmosphere. However, the GHG optimal control problem has the specificity that changes in atmospheric GHG concentrations as well as in climate lead to nonlinear feedbacks in the climate system. In addition, the problem is also stochastic in its nature, as climate has a natural variability and as emissions of GHG can only be imperfectly controlled by policy makers. To tackle these issues, I develop a stochastic optimal pollution management model where the stock pollutant dynamics exhibit a non-convex feedback.[1] I calibrate the model to represent carbon dioxide (CO_2) accumulation in the atmosphere and the mitigation of climate change.

Accounting properly for feedbacks in devising an optimal climate policy is crucial. Not only is most of the variation in the estimation of climate sensitivity attributable to differences in feedback modeling in global circulation models (Webb et al., 2006), but the non-linearity of some of these feedbacks makes the climate system even more difficult to predict. Examples of nonlinear feedbacks include the albedo-ice effect (Hansen and Nazarenko, 2004) and the thermohaline circulation (Rahmstorf et al., 2005).[2] Feedbacks between the cycle of carbon emissions and its sequestrations by natural sinks are also increasingly found to be inadequately modelled by linearization. Zickfeld et al. (2011) find that, in a coupled climate-carbon model, land and oceans carbon sinks are less effective than predicted by a linear combination of concentration-carbon and climate-carbon

[1]While non-convex feedbacks are a particular type of non-linearity, it often corresponds to what the climate literature is describing as nonlinear feedbacks. Indeed, nonlinear climate or carbon feedbacks are often assumed to take-off at some point only to stabilize or level-off latter. In any case, the resulting non-linearity can be represented by a non-convex function, usually a convex-concave one.

[2]For a more detailed description of the relationship between the albedo-ice effect, the thermohaline circulation, and climate change see Clark et al. (1999).

cycle feedbacks. Complex ecosystems also have the potential to generate nonlinear responses to climate feedbacks, which traditional carbon-cycle models render poorly (Heimann and Reichstein, 2008).

Nonlinear feedbacks in the climate system can lead to tipping points or thresholds in the climate system. Once a tipping point is reached, the climate system can start to change quickly without anthropogenic emissions needing to change quickly as well. Examples of tipping points include the melting of the Arctic ice sheet, the instability of the West Antarctic ice sheet, losses of permafrost and tundra, Atlantic deep water formation and others (Lenton et al., 2008).

Most economic literature on optimal climate policy has taken one of two approaches. The first focuses on linear feedbacks, which can be summarized in one parameter, referred to as climate sensitivity (Roe and Baker, 2007). This linear formulation does not lead to tipping points or thresholds in the resulting climate model. This is the approach of Ramsey-Cass-A models such as the DICE model (Nordhaus, 2008) and the model of Aaheim (2010). An alternative approach introduces thresholds in the GHG stock which, when crossed, trigger either catastrophic negative payoffs, or a change in the climate regime that exacerbates climate change (Tahvonen and Withagen, 1996). This approach has a stochastic analog that uses a hazard rate to represent the probability of crossing into the catastrophic regime, as in Gjerde et al. (1999) and Tsur and Zemel (2009). In either the deterministic or stochastic setting, most thresholds introduced in climate change management models are exogenous, meaning the threshold, or the probability of crossing it, is independent of the decision maker's actions (Tsur and Zemel, 1996, 1998; Naevdal, 2006, 2007; Fisher and Narain, 2003). Fewer models consider endogenous thresholds that depend on variables controlled be the decision maker (Lemoine and Traeger, 2010).

Nonlinear feedbacks have not yet found their way into economic models of optimal GHG emissions. Moreover, there is debate over how these feedbacks should be incorporated into climate models. Zaliapin and Ghil (2010) critique the one parameter climate sensitivity approach of Roe and Baker (2007) that has become widely used. The disagreement persists as Roe and Baker (2011) have rejected this critique, only to be challenged again by Zaliapin (2011). However, problems with non-convexities have been tackled in other contexts. The seminal work of Skiba (1978), on the optimal growth of a one sector economy with a convex-concave production function, introduced the possibility of decision thresholds, since then referred to as "r points." In environmental economics, similar models were developed later to deal with stock pollutants that exhibited convex-concave dynamics. Tahvonen and Salo (1996) look at a general stock pollutant with a concave-convex decay function, while several articles on the so called "Shallow Lake Problem," including Carpenter et al. (1999) and Brock and Starrett (2003), looked at a similar problem but with a positive convex-concave feedback on the pollution stock. All these models share the conclusions that the optimal steady state is not necessarily unique, and that the control rule is likely non-monotonic and possibly discontinuous. All these conclusions are in striking contrast to the received wisdom on GHG optimal policy, a monotonic and continuous increase in abatement.

The model developed here is the standard stock pollution control problem with a non-convex feedback term and a stochastic term included in the pollution accumulation dynamics. The application to climate change adapts the framework developed by Rezai (2010) to fit an autonomous control problem. The model of Rezai (2010) is itself a modification of c DICE model (Nordhaus, 2008), where most of the climate module of DICE is replaced by a damage function expressed in terms of atmospheric CO_2 concentration. The non-convex feedback introduced is similar to that found in an energy balance model. Such models, presenting temperature as the equilibrium between incoming and outgoing radiations, have been little used in economic analysis of climate change. This is surprising given that they use the simplest representation of nonlinear feedbacks of any climate models. The resulting combined model, formally incorporating stochasticity, yields a non-convex dynamic optimization problem. It contributes to climate change literature by formally incorporating non-convex feedbacks in the dynamics of the stock of GHG. It also contributes to the literature on non-convex pollution control models by setting the problem in a stochastic context and applying it to climate change.

The non-convexity of the model creates thresholds. There are no environmental threshold at a particular point of the stock of CO_2, but the non-convex feedback leads to a threshold region, where the feedback processes become significant. The threshold region in the climatic system is exogenous as it depends only on model parameters. It can lead to an endogenous decision threshold, i.e. a level of CO_2 concentration at which the level of optimal emissions changes discontinuously. The resulting discontinuous control rule determines the probabilities of entering different basins of attraction.

The rest of this paper is organized as follows. The generic model is presented in section 2. Section 3 presents the functional forms used to represent the problem of optimal CO_2 emissions and presents some analytical scenarios. The numerical approach used to solve the model, the calibration, and the results are presented in section 4. Section 5 concludes.

2 Model

The general problem involves a productive activity yielding utility that also generates pollution, which accumulates into a stock. This stock reduces utility either directly or indirectly by negatively affecting the productive activity.

This pollution t model is fairly standard in environmental economics. However, if the dynamics of the stock pollutant display non-convex feedback, as in the case of climate change, the problem becomes much more complex and much less studied (Tahvonen and Salo, 1996; Carpenter et al., 1999; Brock and Starrett, 2003). This paper contributes to that literature by setting the problem in a stochastic context and applying it to climate change.

The problem can be described as the maximization of the expected sum of discounted net benefits from the productive activity, subject to the dynamics of the pollution stock. Formally:

$$\max_{\{x_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} \mathbf{E}_{\xi} u(x_t, S_t, \xi_t) \tag{1}$$

subject to

$$S_{t+1} = f(x_t, S_t, \xi_t) \tag{2}$$

$$x_t \in \Phi(S_t) \tag{3}$$

where:

- $x_t \in \mathcal{X} \subseteq \mathbb{R}$ is the productive activity (control variable) at time t;
- $S_t \in \mathcal{S} \subseteq \mathbb{R}$ is the stock of pollution (state variable) at time t;
- ρ is the positive discount rate;
- $\xi_t \in \Xi \subset \mathbb{R}$ is i.i.d. and Ξ is finite;
- $u : \mathcal{X} \times \mathcal{S} \times \Xi \to \mathbb{R}$ is the net benefit from productive activity x_t, pollution stock S_t and stochastic shock ξ_t at time t, where $\frac{\partial u}{\partial x} > 0, \frac{\partial^2 u}{\partial x^2} < 0$, $\frac{\partial u}{\partial S} < 0$, $\frac{\partial^2 u}{\partial S^2} < 0$;
- $f : \mathcal{X} \times \mathcal{S} \times \Xi \to \mathcal{S}$ is the transition function for the stock;
- $\Phi : \mathcal{S} \to 2^{\mathcal{X}}$ is the feasible action correspondence.

The existence of a solution to this problem depends on conditions that can be imposed on the one period reward function, the discount rate, the transition function and the feasible action correspondence. If $\mathbf{E}_\xi u(x_t, S_t, \xi_t)$ is real valued, continuous and bounded, $\rho > 0$, f is continuous and Φ is non-empty, compact-valued and continuous, then there exist a unique value function $V(S_t, \xi_t)$ that solves:[3]

$$V(S_t, \xi_t) = \max_{x_t} \left\{ u(x_t, S_t, \xi_t) + \frac{1}{1+\rho} \mathbf{E}_\xi V(S_{t+1}, \xi_{t+1}) \right\} \tag{4}$$

None of the stated restrictions on the problem preclude the use of a non-convex feedback term in the pollution stock dynamics; the only additional restriction needed is that the overall behaviour of this stock be continuous. The next section

[3]The proof of that theorem is that of proposition 2 in chapter 6 of Bertsekas (1976).

shows that while the dynamics of the CO_2 stock is continuous, the location of steady states to the system can change discontinuously in the level of emissions.

Because of the infinite time horizon, the fact that the reward function does not depend on calendar time, and the distribution of ξ_t is time invariant, this is an autonomous dynamic programming problem. Hence the value function does not have time as an explicit argument. Such an autonomous problem has an optimal policy function that is stationary.

The first order condition of equation (4) gives a condition relating the marginal net benefit of the productive activity with the marginal cost imposed by the pollution stock along the optimal trajectory.[4]

$$u_x(x_t, S_t, \xi_t) = -\frac{1}{1+\rho}\mathbf{E}_\xi V_S(S_{t+1}, \xi_{t+1}) f_x(x_t, S_t, \xi_t) \tag{5}$$

One interpretation of equation (5) is that for a given x_t, S_t, ξ_t combination to be optimal, it must be the case that the marginal utility derived from the productive activity is equal to the discounted expected marginal reduction in future utility due to extra pollution, weighted by the contribution of that marginal unit of productive activity to the future stock of pollution. Because this condition must hold for all time periods, there is no opportunity for arbitrage. Hence, on an optimal path, it is impossible to rearrange the timing or total amount of productive activity such that the sum of discounted expected net benefits is increased.

3 Application to Climate Change

The application to climate change that I propose builds upon a modification of the DICE model[5] developed by Rezai (2010). While Rezai (2010) simplifies the climate module of the DICE model such that all stock dynamics and climate damages can be represented by a single stock variable (CO_2 concentration in the

[4] I use the notation u_x to denote the partial derivative of the function u with respect to variable x.

[5] The version of the model I will be referring to is DICE-07 (Nordhaus, 2008).

atmosphere), I also simplify the economy side of the model by assuming constant capital stock, constant savings, and constant production and abatement technology. These simplifications allow me to keep the model analytically and numerically tractable when I introduce non-convex feedback and uncertainty.

In the next subsection, I present the functional forms of the climate application of the model and discuss in more details how these depart from the original DICE model and its modification by Rezai (2010). In the subsequent subsections, I present some further analytical results to establish what type of scenarios are possible given the choice of functional forms.

3.1 Applied model

While many models of climate change mitigation, such as DICE, are optimal growth models with exogenous technical change, I make the problem static in these two dimensions. This modification leads to an autonomous problem, which is more tractable for analytical optimal control and more stable when solving for the numerical control rule. In addition, dropping these exogenous changes does not fundamentally alter the problem at hand. Eliminating output growth reduces future wealth, which in turn reduces the incentive to shift abatement costs to future periods. In addition, on the technological side, eliminating the decline in abatement costs also reduces the incentives to shift abatement to the future. Part of these effects are counterbalanced by the constant business as usual level of emissions, which is due to the constant output and emissions intensity assumptions. Indeed since emissions are not expected to grow over time, there is less pressure to act now than in the standard DICE model.

The model presented in the previous section can be parametrized using the following functional forms to represent the mitigation of GHG emissions. The one period payoff, or reward function, is a utility function with constant elasticity of marginal utility, $u(m_t) = \frac{x(m_t)^{1-\eta}}{1-\eta}$. The choice variable, $m_t \in [0, 1]$, is the share of output devoted to the mitigation of CO_2 emissions. This variable completely determines per i consumption, $x(m_t)$, according to the function $x(m_t) = \frac{\bar{y}d_t(1-m_t-\sigma)}{N}$, where $\bar{y} > 0$ is the constant gross world output,[6] $d_t \in [0, 1]$ is the fraction of gross output remaining after climate damages have been accounted for, $\sigma \in [0, 1)$ is the constant savings rate, and $N > 0$ is the constant population level.

To be able to numerically solve the model, given the potential for a c point in a stochastic setting, the number of state variables must be kept to a minimum. Indeed the curse of dimensionality is quite acute in non-convex problems, as the concept of Skiba points must be expanded to Skiba lines, planes or hyperplanes, as the number of state variables grows to two, three or more. To avoid these complications, I limit the model to one state variable, which is the stock of CO_2 in the atmosphere, S_t, measured in parts per million volume (ppmv).[7]

Since CO_2 is the only state variable in the model, the climate damages must be expressed in terms of this stock. I use the formulation developed by Rezai (2010), where the fraction of output spared from climate damages is defined as:

$$d_t = \left(1 - \left(\frac{S_t - 280}{\bar{S} - 280}\right)^{1/\gamma}\right)^{\gamma}. \tag{6}$$

This damage function can lead to more or less damages than in the DICE model, depending on the value of $\gamma \in (0, 1)$. This parameter defines the curvature of the damage function, where damages get linear as $\gamma \to 1$. $\bar{S} > 280$ is the stock of CO_2, measured in ppmv, that leads to the complete destruction of world output.

I also adapted the formulation of the transition function from Rezai (2010). Indeed, his recast-DICE model simplifies the climate module of DICE by replacing it by a transition function that depends only on the stock of CO_2. This approach bypasses most of the climate module of DICE, which eliminates several intermediate state variables such as carbon in upper and lower oceans, and atmospheric and oceanic temperatures. Such a streamlined formulation simplifies the analysis and lends itself more readily to numerical computations. The drawback is that some of the delays in the climate system are lost and the damage function must be specified in terms of CO_2 concentration instead of temperature.

[6] Measured in trillions of dollars.

[7] The pre-industrial concentration of CO_2 is fixed at 280 ppmv (Randall et al., 2007).

The transition function, f, describing the evolution of the stock of CO_2 between each period is described by equation (7).

$$f(m_t, S_t, \xi_t) = (1-\varepsilon)S_t + (\beta - M_t)y_t + g(S_t) + 280\varepsilon + \xi_t \tag{7}$$

Hence the stock of CO_2 next period depends on the current stock, S_t, which dissipates at rate $\varepsilon \in (0,1)$, on the level of output net of climate damages, $y_t = \bar{y}d_t$, on a non-convex feedback term, $g(S_t)$, and on a stochastic term, ξ_t. Output has a constant CO_2 intensity, $\beta > 0$, which can be reduced through mitigation, M_t, defined as the per unit of output reduction in CO_2 emissions due to mitigation m_t. The relationship between M_t and m_t is represented by the function $M_t = \left(\frac{m_t}{\theta_1}\right)^{1/\theta_2}$, where θ_1 and θ_2 are positive parameters that affect the output cost of a given reduction in emission intensity as defined in DICE.

The new feature in this climate model is the feedback term, $g(S_t)$, in the transition function for the stock of CO_2. The functional form chosen for $g(S_t)$ is given in equation (8).[8]

$$g(S_t) = \mu(\tanh(\kappa(S_t - \hat{S})) + 1) \tag{8}$$

This function is convex-concave, where $\mu > 0$ is a scaling parameter of the feedback function such that $\max g(S_t) = 2\mu$, $\kappa > 0$ is a parameter affecting the slope of the feedback function, and $\hat{S} > 280$ is the stock of CO_2 at which the feedback function reaches its inflexion point. As will be illustrated in the next section with specific values, one can describe μ as the magnitude of the feedback, κ as the suddenness of the feedback, and $\hat{S}$ as the midpoint of the region where the feedback takes off, or the onset region.

Defining this feedback term as an increasing convex-concave function can represent several climate scenarios. In one of its simplest representations, the climate system can be modelled as a zero dimensional energy balance system.[9] The

[8] This formulation is similar to that of Zaliapin and Ghil (2010), who use it in the context of an energy balance model.

[9] The dimensionality of energy balance models refers to their level of spatial aggregation.

planet's temperature is the resulting equilibrium of the incoming energy absorbed from solar radiation and the outgoing energy emitted by the earth. As the temperature changes, the ice cover on the planet changes, which in turn changes the amount of solar radiation absorbed by the system, as more or less ice will reflect more or less radiation. This feedback effect makes the relationship between incoming energy and temperature non-convex.

Despite the model developed in this paper not being one of energy balance, a similar reasoning can be applied to a model of CO_2 concentration balance. As CO_2 concentration rises, temperature also rises, which sets in motion changes in carbon sources and sinks that will further increase the concentration of carbon. Such positive carbon cycle feedbacks have been found using recent data (Cox et al., 2000; Heimann and Reichstein, 2008; Randall et al., 2007) and also with statistical studies based on paleoclimatic data (Lemoine, 2010). The evidence is mounting that these feedbacks are nonlinear and lead to non-convexities in the climate system (Zickfeld et al., 2011; Friedlingstein et al., 2001). Since the model developed here is highly aggregated, like zero dimensional energy balance models, the specific functional form of $g(S_t)$ does not represent a specific climate process, but, as in Zaliapin and Ghil (2010), it is a smoothed version of the more usual step or piecewise linear functions used to describe non-convex dynamics.[10]

Finally, the last innovation in the model is the presence of uncertainty in the form of an additive random term ξ_t. By making the next period stock of CO_2, and therefore the damages, uncertain the model encapsulate risks of crossing thresholds, beyond which the carbon cycle will change significantly. Managing these risks is an important part of developing comprehensive policies to address the challenges posed by climate change.

Zero dimensional models consider the planet surface temperature as uniform, one dimensional models distinguish temperature based on latitude, and two dimensional models distinguish both latitude and longitude.

[10]In energy balance models, the piecewise linear formulation is referred to as Sellers-type models (Ghil, 1976), while the piecewise constant formulation, or step function, is referred to as Budyko-type models (Held and Suarez, 1974).

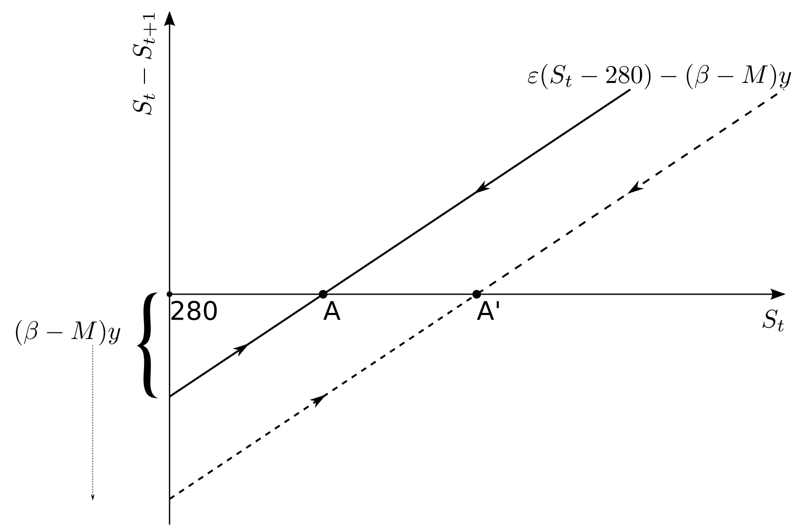

Figure 1: Linear CO_2 dynamics with constant emissions

3.2 Climatic thresholds in a deterministic setting

Thresholds are best introduced by first tackling the deterministic analog of the model, where ξ_t is always zero. In this subsection, as well as the subsequent ones, I use phase diagrams to represent the characteristics of different formulations of the model. In general, phase diagrams might not accurately represent the dynamics of a discrete time model, as discrete shifts in state and control variables can make the system jump around a point that would otherwise be monotonically reached. Appendix A shows that, in the deterministic version of the model, the stock of CO_2 is always monotonically approaching the stable steady states. Hence, phase diagrams provide an accurate rendition of the model dynamics.

On the purely climatic side, one must first understand the behavior of the stock of CO_2 for constant levels of emissions before moving to the optimally controlled climate-economy. Figure 1 presents this situation when there is no feedback. With constant emissions $(\beta - M)y$, the long run stock steady state is point A. If the initial stock, S_0, is below A, the CO_2 concentration will increase to asymptotically reach point A, while it would decrease to asymptotically reach point A, if S_0 is above A. If M permanently decreases, inducing an increase in emissions, the long

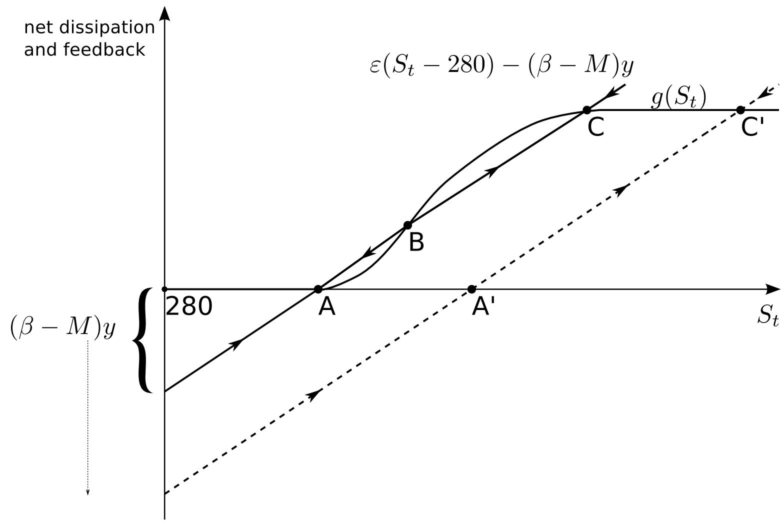

Figure 2: Nonlinear CO_2 dynamics with constant emissions

run steady state shifts proportionally to point A'.

In the nonlinear case, figure 2 presents the deterministic components of f. Intersections between the straight dissipation line and $g(S_t)$ now represent the steady states of the climate system. I will refer to the part of $g(S_t)$ to the left of its inflection point as the lower branch and the part to the right of that point as the upper branch. Points A and C are locally stable equilibria, while point B is an unstable one. Hence, for constant emissions $(\beta - M)y$, if S_0 is less than B, the stock of CO_2 would asymptotically converge to point A, while if S_0 is greater than B, the stock of CO_2 would converge to point C. As in the linear case, consider a permanent decrease in M. If that decrease is large enough, such that the shifted dissipation line now intersects only once with $g(S_t)$, only one steady state remains, C', and it is now globally stable.

It is now easy to construct an example where a permanent change in emissions leads to a more than proportional change in the steady state stock of CO_2. Suppose that S_0 is less than B and emissions are $(\beta - M)y$. If M permanently decreases such that emissions increase as shown in figure 2, the steady state stock

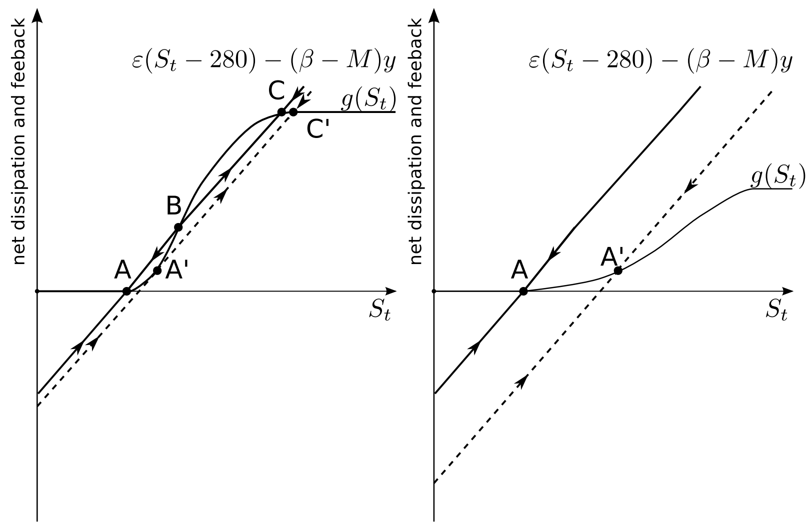

Figure 3: Nonlinear CO_2 dynamics with and without hysteresis

of CO_2 will not shift to A' as in the linear case, but to a much higher C'.[11] That is, the increase in emissions and their subsequent accumulation in the atmosphere have triggered feedbacks in the climate-carbon cycle such that the ultimate steady state of the CO_2 is much higher than it would have been without these feedbacks.

To understand how thresholds come to be in the climate system, an important distinction must be made between two potential cases that the functional form of $g(S_t)$ can represent. Those are depicted in figure 3. In the left panel, the maximum slope of $g(S_t)$ is greater than ε, while in the right panel it is smaller. In the left panel, there are two thresholds in the stock of CO_2. Starting from a stock of CO_2 below A, as emissions increase, the steady state stock moves along the lower branch of $g(S_t)$ up until point A', where the dissipation line and $g(S_t)$ are tangent. If emissions go beyond the intercept of the doted line on the left panel of figure 3, then the steady state level of emissions would go beyond point C' on the upper branch of $g(S_t)$. In addition, if emissions would decrease back to the level where the steady state was previously A', the steady state would not go back to A' but would stay at C'. This asymmetry in the relationship between

[11]Note that exactly the reverse example can be created if S_0 is above B and M increases enough.

emissions and stock steady state is referred to as hysteresis. Hence, the stock value corresponding to point A' is a threshold, below which the steady state moves along the lower branch of $g(S_t)$. There is a analogous point to A', between B and C along $g(S_t)$, beyond which the steady state moves along the upper branch of $g(S_t)$. It is the second tangency point between the dissipation line and $g(S_t)$.

In the case shown in the right panel, no thresholds exist, as $g(S_t)$ is nowhere steeper than the dissipation line. As emissions increase, the stock steady state also increases along $g(S_t)$. Also, the absence of thresholds imply there is no hysteresis effect. Indeed, the steady state A' in the right panel of figure 3 can be reached from above or below monotonically.

I focus on the hysteretic case depicted by the left panel of figure 3. This case is likely the most accurate representation of the management of CO_2 emissions. Indeed, while ε is very small because carbon sinks retire CO_2 relatively slowly from the atmosphere, the maximum slope of $g(S_t)$ is likely to be very high, because nonlinear feedbacks are believed to kick in over a relatively short range of stock values. Budyko-type models use a step function to model such feedbacks, which leads to an infinite slope at the point of discontinuity.

3.3 Climatic thresholds in the stochastic setting

The introduction of the stochastic term into the function f changes the interpretation that can be given to climatic thresholds. To tackle this change, I build on the previous framework where emissions were kept constant over time, instead of being adjusted to follow an optimal policy. The behaviour of the problem with an optimized emissions path is addressed in the next subsection.

Because of the randomness of ξ_t, one must now think of steady state distributions instead of just steady states. The steady state distribution can be represented fairly easily in the linear case. Adding a vertical axis to the right of figure 1 to represent the probability density function of the steady state, this steady state distribution can be represented as in figure 4. The situation depicted assumes that ξ_t has expectation zero. In this case, A is the mean of the steady

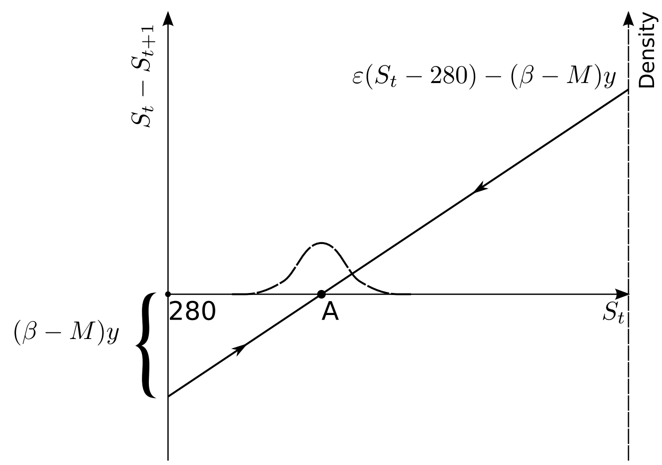

Figure 4: Linear stochastic CO_2 dynamics with constant emissions

state distribution instead of being the steady state itself.

A similar exercise can be done with the nonlinear dynamics. Figure 5 shows how the density of the steady state distribution could look under two scenarios. Both scenarios assume that S_0 is below B, such that the deterministic steady state is point A. In the stochastic case, if the variance of the random term is low relative to the dissipation rate, the steady state distribution is clustered around A.[12] That is, no sequence of random shocks is strong enough to push the long run CO_2 stock towards C. In contrast, if the variance of the stochastic term is higher, then it is possible that a sequence of shocks pushes the CO_2 stock beyond B, where the deterministic part of the stock dynamics could keep pushing the stock in the direction of C. That is, even when S_0 is below B, there would be a non-zero probability that the long run stock of CO_2 is in the neighbourhood of C. In such a case, the steady state distribution is bimodal, as shown in figure 5.[13]

As the construction of these scenarios show, it is possible that by starting on

[12] Note that the distribution is not symmetric anymore because the deterministic part of the stock dynamics is not symmetric.

[13] There is a third possibility that is not depicted in figure 5. The variance could be so big that the stochastic part of the stock dynamics completely overwhelms the deterministic part, such that the density of the steady state distribution is single peaked and encompasses both A and C. I however disregard this possibility as uninteresting, because it means that the system is so volatile that it can barely be controlled.

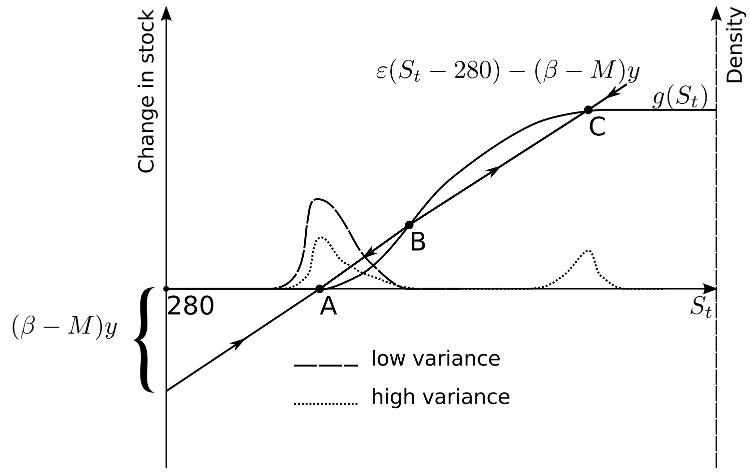

Figure 5: Nonlinear stochastic CO_2 dynamics with constant emissions

either side of point B, the system would reach a different steady state distribution, gravitating towards either point A or point C. That is, the steady state distribution might not be unique. If it is non-unique, new thresholds arise in the model. With constant emissions and a deterministic stock, the dynamics of the model were quite simple. Start to the left of B, end up at A and start to the right of B, end up at C. Now with the stochastic term, there is some uncertainty when S_0 is in a certain neighbourhood of B as to the direction that the long run stock of CO_2 will take. However, if we are in the low variance case described before, it is possible that beyond certain stock values, away from B, it might be certain that the stock will not go back beyond B. Figure 6 illustrates such a situation. In this figure, D and E represent threshold stock values that mark the limit of a basin of attraction, i.e. a region that once the stock has entered, it will never leave. When the CO_2 concentration is between D and E, there is a positive probability that it will enter either of the basins of attraction. In the context of climate policy, these thresholds might be perceived as bounding risk regions. For example, with this constant emissions policy, the decision maker would know that if the stock gets above point D, there is a risk of ending up with a much higher CO_2 concentration, even if emissions remain unchanged.

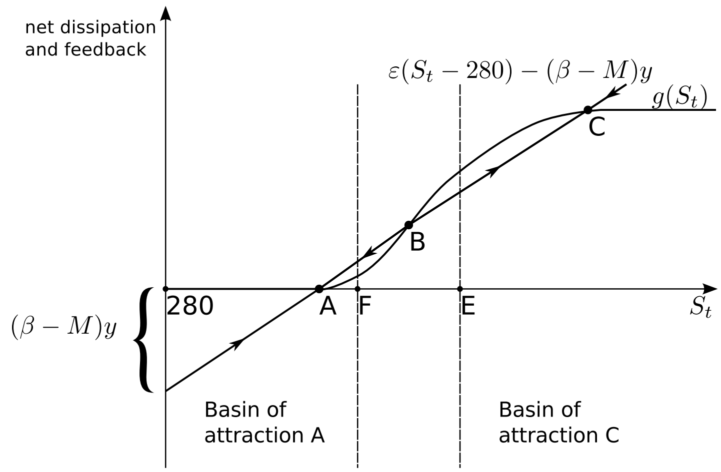

Figure 6: Basins of attraction for nonlinear stochastic CO_2 dynamics with constant emissions

3.4 Decision thresholds

Dropping the constant emissions assumption, I analyze the optimally controlled climate-economy. The optimal control of CO_2 emissions involves a policy rule where the level of emissions varies according to the observed stock of CO_2. As with the study of stock dynamics, it is useful to first look at a deterministic version of the model to highlight a few of its specificities before moving to the stochastic analog. The first step is to derive the two isoclines in the model, the stock and the mitigation isoclines. The stock isocline derivation requires only to equate f with S and then solve for M/β, the fraction of emissions abated.[14] Hence the stock isocline is:

$$\frac{M}{\beta} = \frac{(280 - S)\varepsilon + g(S)}{\beta \bar{y} d} + 1 \tag{9}$$

[14] All the phase diagrams will be presented with M/β, the fraction of emissions abated, as a function of S. Technically M/β is not the choice variable, but it is a convenient transformation that facilitates interpretation.

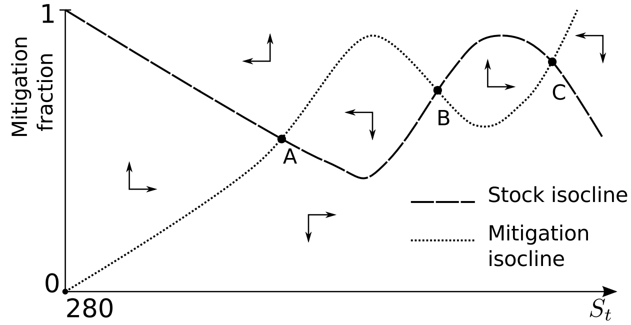

Figure 7: Phase diagram of the deterministic climate-economy

To find the mitigation isocline, I use the first order condition. I rewrite the first order condition in equation (5) such that it represents the deterministic problem.

$$u_m(m_t, S_t) = -\frac{1}{1+\rho} V_S(S_{t+1}) f_m(m_t, S_t) \tag{10}$$

I apply the envelope theorem to the deterministic value function to get a second marginal condition.

$$V_S(S_t) = u_S(m_t, S_t) + \frac{1}{1+\rho} V_S(S_{t+1}) f_S(m_t, S_t) \tag{11}$$

Combining equations (10) and (11) and dropping the time subscripts yields the condition for the mitigation isocline.

$$u_m = \frac{1}{1+\rho} \left[f_S u_m - f_m u_S \right] \tag{12}$$

Substituting the appropriate partial derivatives by their parametric expressions, I obtain an implicit expression for the mitigation isocline, equation (13). This equation can be numerically solved for given parameter values.

$$\frac{1}{\theta_1 \theta_2}(1-\sigma)\left(\frac{M}{\beta}\right)^{1-\theta_2} + \left(1 - \frac{1}{\theta_2}\right)\frac{M}{\beta} = \frac{g_S - \varepsilon - \rho}{\beta \bar{y} d_S} + 1 \tag{13}$$

Both isoclines can be represented in a phase diagram like figure 7. It illus-

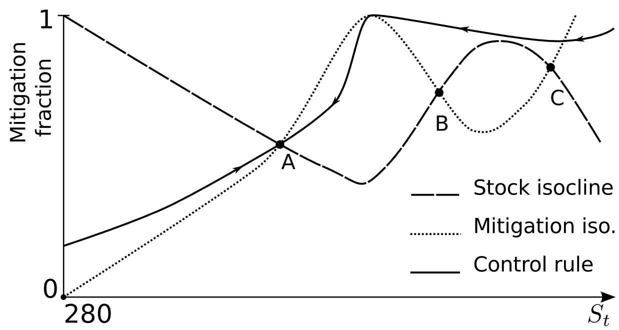

Figure 8: Control rule 1 (deterministic climate-economy)

trates a situation where the optimally controlled climate-economy has three steady states; points A and C are stable, while B is unstable. In each sector delimited by the isoclines and the axes, directional arrows show qualitatively how the stock of CO_2 and the fraction of emissions abated evolve over time. This is only one possibility for the number of steady states. It could be the case that the isoclines only cross one time, either early at a point like A, or late at a point like C. In such cases, there would be only one steady state. I focus on the depicted case where there are three steady states, both because it is the most frequent in the numerical implementation and because it conceptually encompasses the other cases.

Using the three steady states phase diagram, I show there are three qualitatively different possibilities for the control rule. The first possibility is having a control rule that leads steady state A to be globally optimal. This situation is depicted in figure 8. In this case, regardless of the value of S_0, the long run level of mitigation and CO_2 in the atmosphere will be point A. Note that such a control rule has to be non-monotonic in abatement, because for stock values between B and C, it must go through a region of the phase diagram where mitigation is increasing, while mitigation is decreasing in the two regions adjacent to the middle one. Conversely, it could be the case that steady state C is globally optimal, resulting in a control rule that looks like that shown in figure 9. Once again, for a reasoning similar to that of the previous case, the control rule has to be non-monotonic in CO_2 concentration.

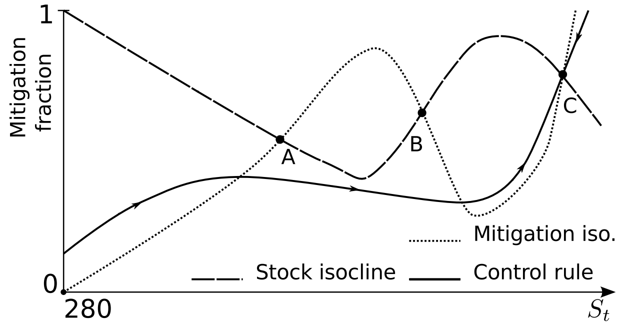

Figure 9: Control rule 2 (deterministic climate-economy)

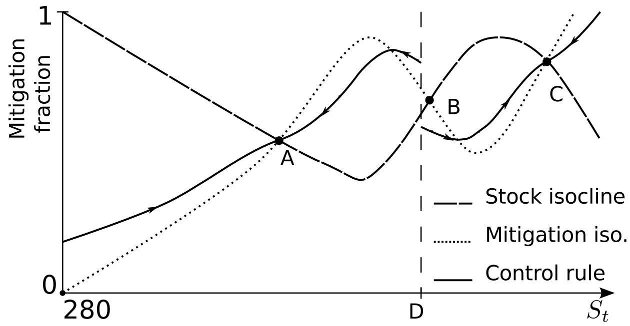

Figure 10: Control rule 3 (deterministic climate-economy)

A third possibility can arise where both steady states A and C are locally optimal. This case is shown in figure 10. Here, the control rule is not only non-monotonic in abatement, but it is also discontinuous. The point at which the discontinuity occurs, point D, is the Skiba point. The Skiba point is a decision threshold. When S_0 is arbitrarily close to D, but smaller than it, it is optimal to abate a relatively large fraction of emissions in order to push the stock of CO_2 towards A in the long run. In contrast, when S_0 is arbitrarily close to D, but larger than it, it is optimal to abate a relatively smaller fraction of emissions and let the stock of CO_2 grow to its steady state value of C. That intuition confirms the mathematical property that at the Skiba point itself, the value function of either trajectory is the same, meaning that the decision maker is indifferent between following the mitigation path that leads to point A or the one that leads to point C.

Just from the shape of the isoclines, one can sometimes rule out certain possibilities. For example, the particular isoclines depicted in figure 8 do not allow for a control rule that makes point C the globally optimal steady state. Point C cannot be approached from an initial stock value below A, as the only phase region that allows the stock to grow from point A to point B also requires the mitigation fraction to decline. When the stock would reach B, the mitigation fraction would be too low to let the system gravitate towards point C. Conversely, the example presented in figure 9 rules out the global optimality of steady state A. Using a similar reasoning as in the previous case, point A could not be reached from an initial stock value above C, as the mitigation fraction between point C and point B would become too high to enter the phase region leading to point A from above. However, the case with a Skiba point can never be ruled out just by the shape of the isoclines, whenever the optimally controlled climate-economy has three steady states.

Moving from the deterministic optimally controlled climate-economy to the stochastic one is similar to the previous discussion of the stochastic climate dynamics with constant emissions. To illustrate the impact of the addition of uncertainty to the deterministic optimally controlled climate-economy, I use the last example where the two stable steady states were optimal in the deterministic case.

Figure 11 is very similar to figure 10 with the addition of probability density functions for the steady state distribution. The control rule, however, no longer intersects the deterministic steady states. This phase diagram keeps the deterministic isoclines only as a reference point for the new stochastic control rule. While the mitigation isocline does not retain a similar interpretation in the stochastic setting, the stock isocline is still meaningful in the stochastic setting as it represents the locus of points where the stock of CO_2 is constant in expectation. Hence one would expect the modes of the probability density function of the steady state distribution to gravitate around the intersections of the deterministic stock isocline and the stochastic control rule. Again, this steady state distribution may or may not be unique and there may be one or more basin of attractions over the state space.

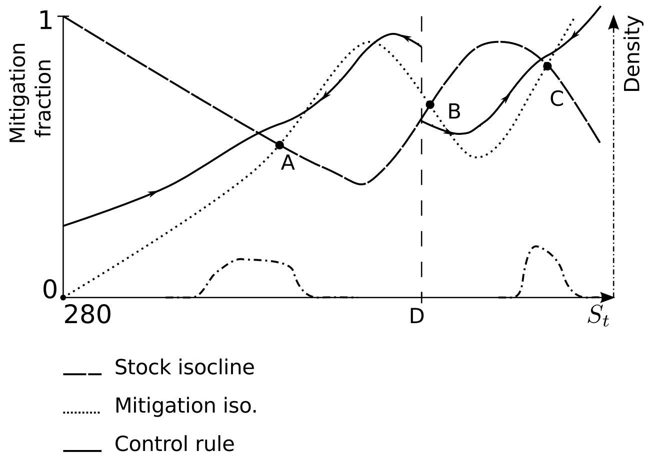

Figure 11: Optimally controlled stochastic climate-economy

The stochastic control rule depicted in figure 11 is arbitrary. There is very little analytical information that can be derived from the problem statement. In particular, the existence and location of a Skiba point cannot be analytically derived. Hence, to explore the influence of different parameters on the control rule and to quantify the climate policy implications of the present model, numerical solutions are necessary. The next section discusses numerical simulations performed with several calibrations of the model.

4 Numerical Implementation

In this section I use numerical techniques to solve for the optimal control rule to the problem posed in the previous section.[15] In order for the results of this numerical exercise to be of interest, I calibrate the parameters of the model to represent the problem of climate change mitigation, within the limits of the current framework. The section begins with a description of the algorithm used to solve for the value function. Then, I present the calibration of reference, carefully explaining where

[15]This implementation draws from ideas presented in Miranda and Fackler (2004) and some of the subroutines contained in the associated toolbox, as well as from notes from 2010 summer program from the Institute for Computational Economics at the University of Chicago.

the value of each parameter comes from. Finally, I do some comparative dynamics and robustness checks by varying the value of some of the key parameters.

4.1 Solution approach

To solve for the optimal control rule of the stochastic climate-economy, I use the collocation method. The possibility of a Skiba point complicates the search for this optimal control rule. Indeed, when such a point arises, the control rule is discontinuous at that point in the state space, while the value function is kinked at that same state location. Because it is easier to approximate kinked but continuous functions than discontinuous ones, I apply the collocation method by using the value function iteration. Once the value function has been solved, it is trivial to back out the control rule.

The collocation method approximates a function, here the value function, as the linear combination of basis functions over nodes forming a finite subset of the state space. If the number of nodes is equal to the degree of the basis functions, the value function V, evaluated at the collocation nodes S_i, is exactly equal to the linear combination of the basis functions ϕ_j at these same collocation nodes:

$$V(S_i) = \sum_{j=1}^{n} c_j \phi_j(S_i) \quad \forall\, S_i \in \{S_1, S_2, \ldots, S_n\} \tag{14}$$

where c_j represents each of the n coefficients of the linear combination that are to be solved for. Starting with an initial guess for the values of the coefficients, those are updated by using the equation

$$c = \Phi^{-1} v(c) \tag{15}$$

where $\Phi_{ij} = \phi_j(S_i)$ and $v_i(c) = \max_m \left\{ U(m, S_i) + \frac{1}{1+\rho} \mathbf{E}_\xi \sum_{j=1}^{n} c_j \phi_j(f(S_i)) \right\}$. That is, one performs the maximization implied by the value function at each of the evaluation nodes to find an approximate value function from which one can recover the implied coefficients. Those new coefficients are then compared to those of the previous round (or the original guess if this is the first round) to decide if the iteration process should be stopped. If the coefficient values derived from equation (15) are close enough to those of the previous round, the process is stopped, otherwise it gets repeated.

To address the particular challenges of approximating a kinked value function, I implement the collocation method with the following specification:

- cubic splines for basis functions;
- modified Chebyshev nodes;
- L_∞ norm to evaluate the distance between two sets of coefficients.[16]

Using splines to approximate the value function, repeated break points at some state node can be used to reduce the smoothness requirements of the spline approximation. Stacking n break points at a given node reduces the degree of continuity requirement of the function by $n - 1$. By using cubic splines, having three breakpoints stacked at one point in the state space allows the approximated function to be kinked at that point, while the derivative of the value function must be continuous everywhere else. This approach is well suited to the current problem of approximating a potentially kinked value function. The main problem that remains is finding the location of the kink. Indeed, this location or even the existence of the Skiba point cannot be predicted theoretically. To this end, I develop an algorithm to update the location of the stacked break points in order to numerically find the kink. Chebyshev nodes are useful because they are denser at the boundaries of the state space, where there is less information to approximate the value function. In the current problem however, it is also good to have denser nodes in the vicinity of the kink to better pinpoint its exact location. To do that, I define Chebyshev nodes separately on the intervals on each side of the presumed kink. That way, the nodes are denser both at the boundaries of the state space but also in the presumed neighbourhood of the kink. Finally, the L_∞

[16] The L_∞ norm takes the maximum distance between any two pairs of coefficients as the distance between the two sets.

norm is used because of the potential for a Skiba point. Small changes in the value function can lead to large implied changes in the control rule. As long as any of the coefficients keep changing significantly, the iteration process is allowed to keep going.

As for the algorithm that updates the location of the stacked break, theory says that the kink in the value function must occur at the same point in the state space where the Skiba point occurs in the control rule. If the stacked break points are misplaced, errors in the evaluation of the value function will generate successive spikes in the control rule that do not line up with the stacked break points location. These spikes emerge because I am trying to approximate a non-smooth function with a twice continuously differentiable combination of splines. In such a case, the stacked break points get moved toward the state value where the largest spike on the control rule has occurred, i.e. where the first derivative of the control rule has the greatest magnitude. The iterations are repeated until the kink in the value function and the discontinuity in the control rule occur at the same state node. Of course, one must keep in mind the possibility that there are no such kink and Skiba point in the problem at hand. In that case, putting stacked break point anywhere will not change the approximation and in the absence of discontinuities in the control rule, the algorithm will stop right away.

4.2 Calibration

I now present the values used for the different parameters of the model in what will be the benchmark specification of the numerical results. The time scale of the model is one period for one decade.

Most numbers are either those of the DICE model at the initial point in time or of their equivalent in Rezai (2010). A few have been updated using the *CIA World Fact Book*. The innovation in terms of calibration is the value that is given to the three parameters of the feedback function.

- $\rho = (1+0.015)^{10}-1$ is the per decade pure rate of time preference as defined in DICE (1.5% annual rate);

- $\bar{y} = 74 \times 10$ is ten times the 2010 world GDP in trillion of dollars (CIA, 2011);

- $\sigma = 0.2$ is the fixed savings rate;

- $N = 6.8$ is the 2010 world population in billion of people (CIA, 2011);

- $\bar{S} = 780$ is the CO_2 concentration in the atmosphere in ppmv that leads to the total destruction of world output (it is the number used by Rezai (2010));

- $\gamma = 0.3$ is the elasticity of the damage function as defined in Rezai (2010);

- $\varepsilon = 0.036$ is the per decade dissipation rate of CO_2 (Rezai, 2010);[17]

- $\beta = \frac{0.063}{2.13}$ is the carbon intensity of output in ppmv per trillion dollars as initially defined in DICE;

- $\theta_1 = 0.051$ and $\theta_2 = 2.8$ as initially defined in DICE;

- ξ_t is the quadrature approximation of a random variable distributed $N(0, 0.5)$;[18]

- $\mu = 5.5$ is a feedback function parameter representing the magnitude of the feedbacks;

- $\kappa = 0.04$ is a feedback function parameter representing its maximum steepness;

- $\hat{S} = 560$ is a feedback function parameter setting the location of its inflection point,

The benchmark feedback function and some of its alternatives are represented in figure 12. The most important parameter of the feedback function for the

[17] One cannot give a half life interpretation to the value of ε as carbon dioxide does not decay but cycles in between the atmosphere and terrestrial and oceanic reservoirs. Hence ε represents the net decadal rate at which these reservoirs absorb carbon dioxide from the atmosphere.

[18] This is the approximation of a definite integral by the weighted sum of function values at specified points. I use 5 points for this approximation.

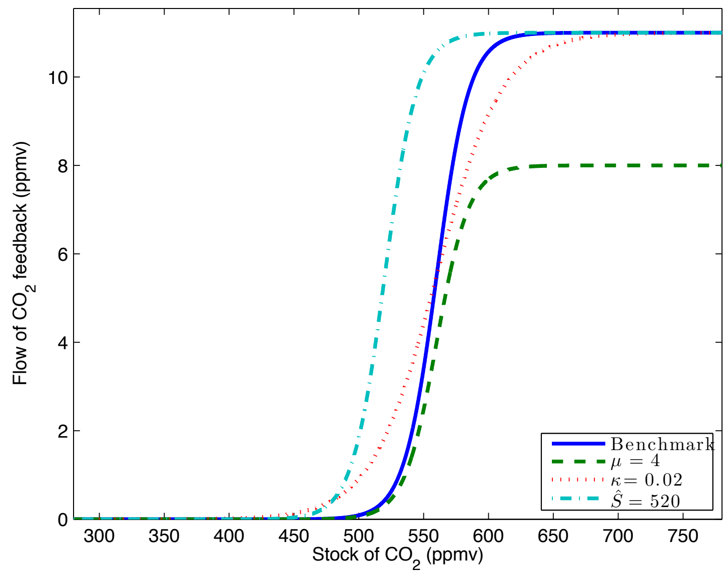

Figure 12: Feedback functions

extent of feedback mechanisms affecting climate change is μ. As it can be seen in figure 12, decreasing μ reduces the magnitude of the feedback in the climate dynamics. κ affects the steepness of the feedback function, a measure of how quickly feedback processes kick in as the concentration of CO_2 in the atmosphere increases. As shown in figure 12, a lower κ leads to a more gradual feedback function, that starts increasing earlier than the benchmark one but that levels off later. Finally, $\hat{S}$, the inflexion point of the feedback function, represents the midpoint of the region where feedback processes take off. By changing $\hat{S}$, one can model feedbacks as arising earlier or later as the CO_2 concentration increases.

As it is most important to correctly calibrate the magnitude of the feedback phenomenon in the climate-carbon cycle, I spend the most time discussing the value of that parameter. While the value of $\hat{S}$ is of quantitative importance, it is of relatively little qualitative importance, as changing its location affects the aggressiveness of the abatement schedule but not its fundamental shape. The case for κ is a little more subtle. As it has already been discussed, the steepness of the

	μ	4			***5.5***			7		
	$\hat{S}$	520	560	600	520	***560***	600	520	560	600
κ	0.02	58.8	56.5	54.2	80.9	77.7	75.6	103	98.9	94.8
	0.04	58.3	56.5	54.6	80.1	77.7	74.5	102	98.8	95.5

Table 1: Difference in year 2100 CO_2 concentration (ppmv) between models with and without feedbacks

feedback function affects the likelihood of the control rule having a Skiba point. More careful robustness checks are therefore required for its value.

I calibrate the magnitude of the feedback function using a meta study of climate-carbon cycle feedbacks. Friedlingstein et al. (2006) apply 11 coupled climate-carbon cycle models to the IPCC[19] A2 emissions scenario to evaluate the increase in atmospheric CO_2 concentration due to feedbacks by 2100. The A2 scenario is one with emissions increasing to very high levels, which is helpful in evaluating the magnitude of the feedback function at very high concentration levels (μ). To select an appropriate value for μ, I simulate the evolution of the CO_2 concentration using the predicted emissions from the A2 scenario and the deterministic climate dynamics (f) with and without the feedback function (g). I report the difference in CO_2 concentrations between the simulations with and without the feedback function for the year 2100 in table 1. What is obvious in this table is that parameters κ and $\hat{S}$ play a relatively insignificant role in the contribution of the feedback function to the stock of accumulated CO_2. This small effect is due to the emissions path used for this simulation being exogenous. These two parameters can affect significantly the control rule and hence the CO_2 stock indirectly. As for μ, it is obviously the main parameter affecting the contribution of the feedback function to the CO_2 stock in the simulations.

The model runs from Friedlingstein et al. (2006) generate feedback contributions to CO_2 concentrations in 2100 between 20 and 200 ppmv. However, most model runs give contributions located between 50 and 100 ppmv. That is why $\mu = 5.5$ is the benchmark value, as it corresponds roughly to the midpoint of

[19]Intergovernmental Panel on Climate Change.

this restricted range. I will also consider values of 4 and 7 for μ as those lead to approximately the lower and upper bounds of that most frequent range.

The final element to calibrate in the model is the stochastic term of the stock dynamics. It represents natural variations in atmospheric CO_2 concentrations. Doney et al. (2006) show that the natural variations in atmospheric CO_2 concentrations over a 1,000 year period has a range of 5 ppm for concentrations around pre-industrial levels and no anthropogenic emissions. To calibrate ξ, I generate 10,000 runs of 100 decades using the carbon dynamics represented by f, without emissions, and a starting value of 280 ppmv. With $\xi \sim N(0, 0.5)$, the median range of variation of the 10,000 runs is 5.44 ppmv and the ninety-fifth percentile is 7.87 ppmv. I hence use this calibration for the benchmark case, as it closely mimics the natural variability of a more complex 3-D global coupled carbon-climate model. However, as pointed out by Joos et al. (1999), data and models suggest that at current higher CO_2 concentrations, the variability is increased.

4.3 Results

I now present the results of the numerical simulation themselves. Before analyzing the benchmark calibration, it is useful to consider what the models yields without the feedback term. Figure 13 shows the control rule and the value function when the problem is solved without the feedback function, g, and the stochastic term ξ. The control rule in this case is a monotonically increasing function of the stock of CO_2, while the value function is a decreasing one. This feedback free formulation highlights that any kink in the value function or discontinuity in the control rule will be attributable to the feedback term and not some idiosyncrasy of the model.

I now turn to the benchmark calibration of the model, including the feedback function and the stochastic term. The control rule and the value function are presented in figure 14. The benchmark calibration yields a discontinuous control rule. The Skiba point lies at 576 ppmv of CO_2. One can see that, as expected, the discontinuity in the control rule is associated with a kink in the value function. Indeed, just before the Skiba point, the value function starts plunging quickly as

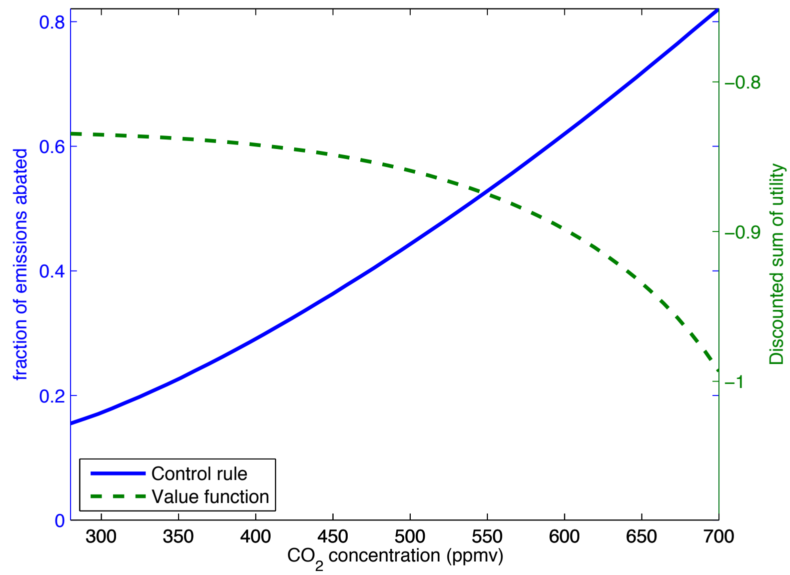

Figure 13: Control rule and value function without feedback and uncertainty

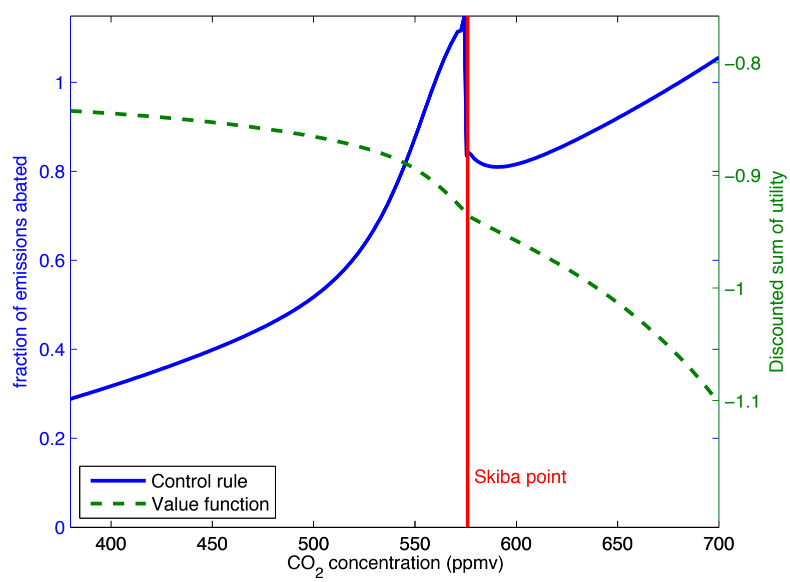

Figure 14: Control rule and value function for the benchmark parametrization

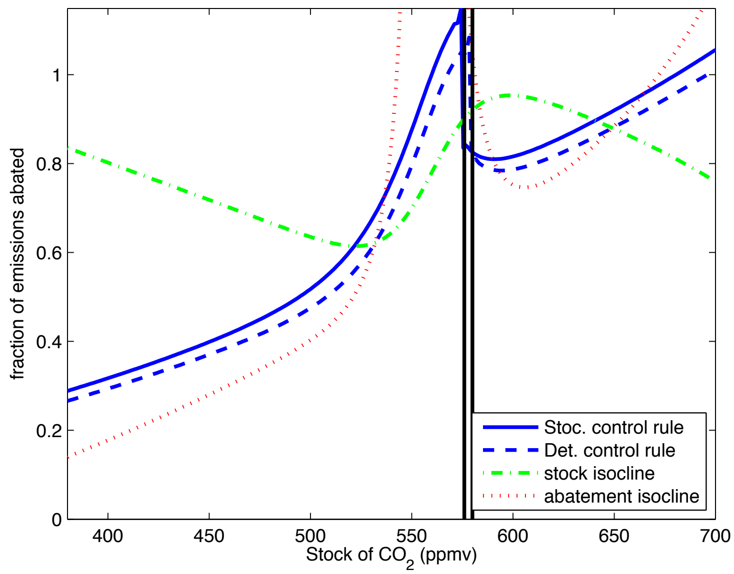

Figure 15: Deterministic vs. stochastic control rules, benchmark parametrization

very high abatement rates, some in excess of 100%,[20] are required to curb the stock of CO_2 to lower concentrations. However, as the stock passes the Skiba point, it is no longer optimal to try to decrease the stock. The sudden reduction in abatement reduces initially the rate at which the value function is falling.

To assess the impact of uncertainty on the control rule, I plot the benchmark control rule and its deterministic analog in figure 15. This figure also includes the isoclines of the deterministic problem. It shows that taking uncertainty into account leads to a more aggressive abatement policy, as the stochastic control rule lies almost everywhere above the deterministic one. The only exception is between the Skiba points of each control rule. Uncertainty lowers the Skiba point, i.e. the threshold at which the stock is allowed to grow towards a higher basin of attraction falls. The intuition for these results is that the decision maker adopts

[20]The model constrains abatement expenses to be less than 100% of GWP, but that does not preclude abatement itself beyond 100%. Such a high level of abatement would imply increasing the carbon sinks beyond their natural levels. That could mean a variety of initiatives, going from reforestation to geo-engineering.

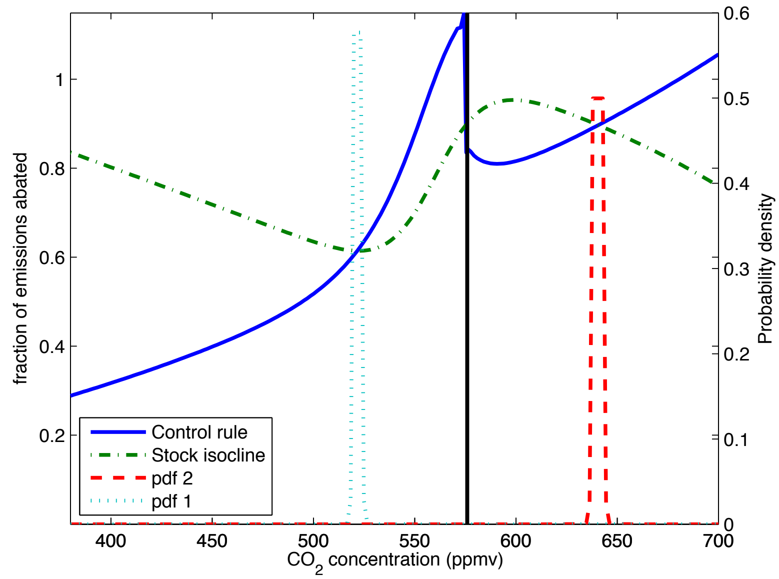

Figure 16: Steady state distributions for the benchmark parametrization

a more aggressive abatement policy to reduce the likelihood of going to higher stocks, because she acknowledges that uncertainty reduces her control on the stock compared to the deterministic case. The extra cost of the more aggressive abatement policy is optimal because it reduces the likelihood of higher stock and hence higher marginal damages, since d is convex. However, the point at which it does not pay to prevent the stock from growing towards the upper basin of attraction is lower, because random shocks might push the stock over the previous threshold in spite of the high abatement effort (hence the lower Skiba point).

The steady state distribution is non-unique in the benchmark case. Using the state nodes as midpoints for bins, I discretize the state space to compute a Markov transition matrix representing the optimally controlled climate economy. The eigenvectors corresponding to the eigenvalues of value one give the steady state distributions of the system. In this benchmark case, there are two eigenvalues equal to one, hence two steady state distributions. The probability density function of each of these is shown in figure 16. Each of the distributions is centred around the intersection of the control rule and the expected stock isocline. The second distribution (higher in stock) is flatter than the first. This difference is not due to changes in the nature of the uncertainty, which is independent of the level of the stock, but to changes in the control rule. Indeed, the control rule being steeper in the neighbourhood of the first steady state distribution means that deviations from the mean of that distribution are corrected faster. For example, if the stock is hit by a negative shock, it will be lower next period but so will the percentage of emissions abated. Hence the stock will be more likely to grow back to its mean steady state value in the subsequent period. Conversely, positive shocks will lead to rapid increases in abatement, which will contribute to curb back the stock. This corrective effect is weaker around the second steady state distribution, since the control rule is flatter there.

Each steady state distribution is contained within a basin of attraction. If the initial stock is far enough from the Skiba point, then starting below it would lead the stock distribution over time to converge towards the first steady state distribution, while it would converge to the second if it started above the Skiba point. Because the uncertainty is fairly small in the benchmark case, as shown by the tight steady state distributions, the space between the basin of attraction is also fairly small. Figure 17 shows the upper bound of the lower basin of attraction and the lower bound of the upper basin of attraction. These bounds are at approximately 574.5 and 577.5 ppmv of CO_2. This interval is centred around the Skiba point (576 ppmv of CO_2). This symmetry is however not represented in probabilities. As shown in figure 18, the probability of ending in the lower basin of attraction decreases very quickly as one moves from the upper bound of that basin to the lower bound of the upper basin of attraction.

Because there is still so much uncertainty about the nature and extent of feedbacks in the climate system, it is important to understand how the control rule changes as the feedback function parameter values are varied. To look at the effect of each parameter, I plot the control rule for alternative parameter values and I compare these to the benchmark. The three panels of figure 19 depict variations in each of the three parameters of $g(S_t)$. In all panels, the solid

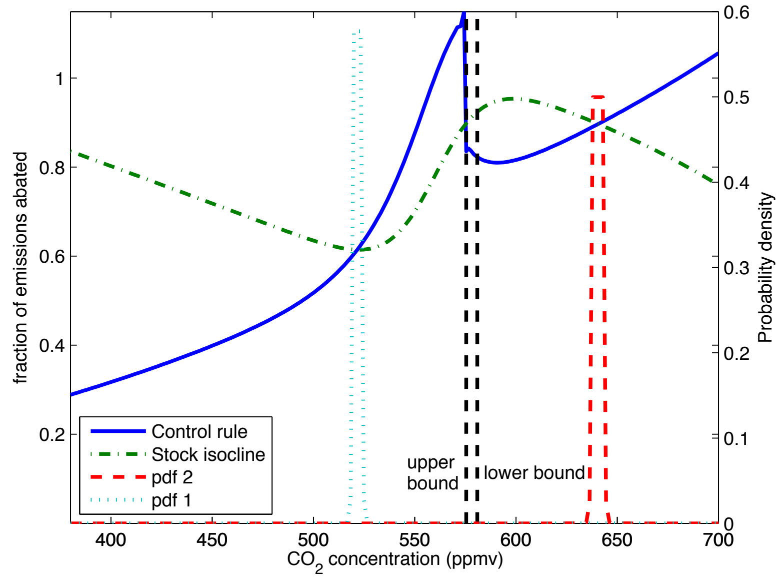

Figure 17: Basins of attraction for the benchmark parametrization

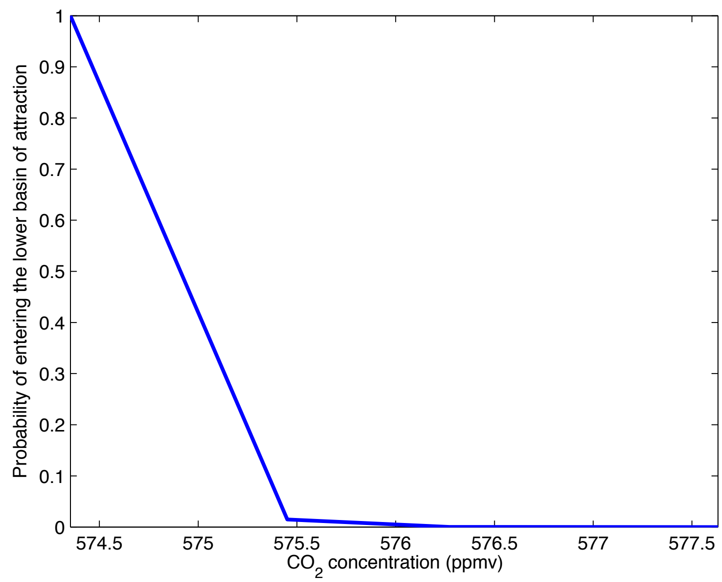

Figure 18: Probability of entering the lower basin of attraction in the inter bound range

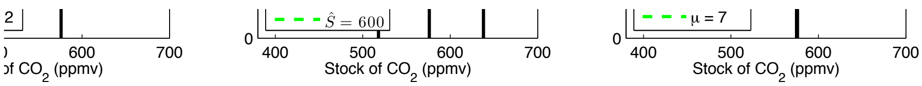

Figure 19: Alternative parameter values for the feedback function

line represents the benchmark control rule. In the left panel, the value of κ is changed such that the steepness of the feedback function is reduced. Reducing the steepness of the feedback function makes the control rule much smoother. In fact, for $\kappa = 0.02$, the control rule is continuous, i.e. it does not have a Skiba point. This distinction means that the suddenness of the feedback onset is crucial to the presence of a discontinuity in the control rule. If feedback occurs moderately gradually (see figure 12 for the shape of the feedback function with $\kappa = 0.02$), then the control rule could become continuous. That is not to say that the steady state distribution in this case would become unique or unimodal. This change in the control rule is the deterministic equivalent of point B in figure 10 becoming an improper unstable node instead of an unstable spiral. In the center panel of figure 19, the location of the feedback[21] is increased and decreased. The primary consequence on the control rule is to change the location of the Skiba point. Earlier onset of feedback implies a lower Skiba point and later onset implies a higher one. This relationship between the onset of feedback and the location of the Skiba point is not surprising as the Skiba point represents a decision threshold that occurs in response to the potential feedback in the stock if that stock is allowed to

[21] Technically, the inflection point of the feedback function.

grow to a certain quantity. What is less obvious is that the location of the Skiba point can be below ($\hat{S} = 520$) or above ($\hat{S} = 560$ and $\hat{S} = 600$) the inflection point of the feedback function. This difference implies that if the feedback sets in earlier, the abatement policy should be very aggressive to prevent the feedback from really kicking in. However, if the feedback begins later, it would be optimal not only to start aggressive abatement later in the stock of CO_2, but later along the feedback curve as well. This asymmetry between the relative position of the Skiba point and the feedback onset is somewhat counterintuitive, as one would have expected the feedback to be more dangerous when it occurs with an already high stock, as damages would be higher. Finally, the right panel in figure 19 shows how the control rule changes as the magnitude of the feedback (parameter μ) increases or decreases. The striking result from this panel is that the feedback magnitude has virtually no impact on the location of the Skiba point. There is, however, an impact on the magnitude of the discontinuity of the control rule at the Skiba point. The higher the magnitude of the feedback, the bigger the discontinuity at the Skiba point. Overall the presence of the Skiba point depends on the suddenness of the onset of feedback, while its location depend on the location of the change in the feedback dynamics. The magnitude of the feedback does not however alter the existence or the location of the Skiba point.

5 Conclusion

The climate system and the carbon cycle are characterized by nonlinear feedbacks. I have modeled these feedbacks as a convex-concave CO_2 feedback function in the stochastic dynamics of the stock of pollution. The impact of this feedback is to create an optimally controlled system with two basins of attraction and a control rule that is potentially discontinuous in between these basins.

The discontinuity of the optimal emissions policy depends on the steepness of the feedback function. A steep function, meaning that the onset of the feedback is sudden, leads to a discontinuous abatement policy function. A flatter feedback function, meaning that the onset of the feedback is more gradual, leads to a continuous abatement policy function. In both cases, the control rule is not monotonic in the stock of CO_2.

If there is a discontinuity in the control rule, its location in the state space depends on the location of the threshold region of the feedback function. If the onset of the feedback occurs at higher CO_2 concentrations, so does the threshold in the optimal abatement policy. The magnitude of this threshold depends on the magnitude of the feedback function. The higher the maximum contribution of feedbacks to the flow of CO_2 into the atmosphere, the higher is the discontinuous change in abatement at the decision threshold.

Stochasticity of the CO_2 stock dynamics reduces the level of the decision threshold. When the natural variability in the stock increases, it is optimal to increase the fraction of emissions abated everywhere, except in a neighborhood of the decision threshold. Since this decision threshold is moved to a lower stock, the fraction of emissions abatement goes down to its right until the previous decision threshold is reached. That is, the smaller the control of the decision maker over the stock, the lower the stock level at which it becomes optimal to let the stock grows toward the higher basin of attraction.

Since the natural variability in the stock of CO_2 is relatively low, the distance between the basins of attraction is small. If this variability were to grow with the concentration of CO_2 in the atmosphere, the distance between the basins of attraction would grow, enlarging the set of stock values where it is uncertain in which basin of attraction the system will end up.

My simulations suggest that when the emissions path is exogenous, the magnitude of the feedback has the largest impact on equilibrium CO_2 concentrations, while suddenness and onset location have comparatively small contributions. This is in contrast to the optimal abatement policy that is most affected by feedback suddenness and onset location, while the magnitude has a comparatively minor impact. It seems that a good share of the empirical climate literature on feedbacks focuses on the magnitude because they consider only exogenous emissions paths. Optimal management of GHG emissions would therefore benefit from more empirical research on the suddenness and onset location of nonlinear feedbacks.

Finally, this framework could be used to tackle the uncertainty in the parameters themselves. Future work could incorporate parameter uncertainty, while also modeling how the decision maker learns about the uncertain parameters over time.

Appendices

Appendix A: Monotonic approach of the steady states

To ensure that the phase diagram exposition is correct in this discrete time setting, I check that the optimally controlled stock approaches stable steady states monotonically. Of course, such reasoning can only apply to the deterministic setting, as random shocks do shift the stock around stable steady states in the stochastic setting.

For the optimally controlled stock to approach the stable steady states monotonically, it must be the case that the function describing the future stock, $f(m_t^*, S_t)$, intersects the 45° line in the plane S_t, S_{t+1} with a positive slope. In the benchmark parametrization, it is the case that $f(m_t^*, S_t)$ has a positive slope everywhere, hence the condition for a monotonic approach of the steady states is satisfied. This is graphically represented in figure 20, which depicts the change in stock, $f(m_t^*, S_t) - S_t$, as a function of the stock. Since the slope of the depicted function is everywhere greater than -1, the stock approaches monotonically each steady state. The evolution of the stock over time is represented by the arrows in figure 20. For any initial stock, one can read the change value on the function itself and report it as the stock next period by drawing a line of slope -1 between the point on the function and the horizontal axis.

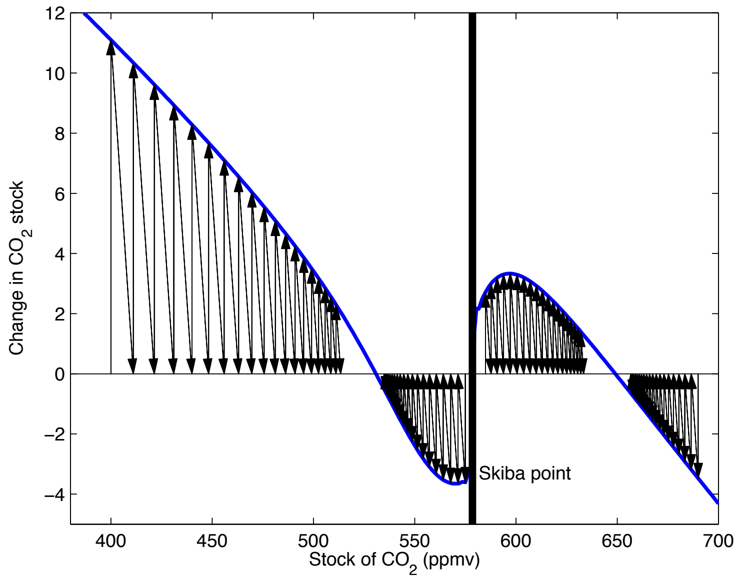

Figure 20: Monotonicity of the steady state approach

References

Aaheim, A., 2010. The determination of optimal climate policy. Ecological Economics 69(3), 562–568.

Bertsekas, D., 1976. Dynamic programming and stochastic control. Academic Pr.

Brock, W., Starrett, D., 2003. Managing systems with non-convex positive feedback. Environmental and Resource Economics 26(4), 575–602.

Carpenter, S., Ludwig, D., Brock, W., 1999. Management of eutrophication for lakes subject to potentially irreversible change. Ecological Applications 9(3), 751–771.

CIA, 2011. CIA World Fact Book. www.cia.gov/library/publications/the-world-factbook/.

Clark, P., Alley, R., Pollard, D., 1999. Northern hemisphere ice-sheet influences on global climate change. Science 286(5442), 1104–1111.

Cox, P. M., Betts, R. A., Jones, C. D., Spall, S. A., Totterdell, I. J., 2000. Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. Nature 408(6809), 184–187.

Doney, S., Lindsay, K., Fung, I., John, J., 2006. Natural variability in a stable, 1000-yr global coupled climate-carbon cycle simulation. Journal of climate 19(13), 3033–3054.

Fisher, A., Narain, U., 2003. Global warming, endogenous risk, and irreversibility. Environmental and Resource Economics 25(4), 395–416.

Friedlingstein, P., Bopp, L., Ciais, P., Dufresne, J., Fairhead, L., LeTreut, H., Monfray, P., Orr, J., 2001. Positive feedback between future climate change and the carbon cycle. Geophysical Research Letters 28(8), 1543–1546.

Friedlingstein, P., Cox, P., Betts, R., Bopp, L., Von Bloh, W., Brovkin, V., Cadule, P., Doney, S., Eby, M., Fung, I., et al., 2006. Climate-carbon cycle feedback analysis: Results from the c4mip model intercomparison. Journal of Climate 19(14), 3337–3353.

Ghil, M., 1976. Climate stability for a sellers-type model. Journal of the Atmospheric Sciences .

Gjerde, J., Grepperud, S., Kverndokk, S., 1999. Optimal climate policy under the possibility of a catastrophe. Resource and Energy Economics 21(3-4), 289–317.

Hansen, J., Nazarenko, L., 2004. Soot climate forcing via snow and ice albedos. Proceedings of the National Academy of Sciences of the United States of America 101(2), 423–428.

Heimann, M., Reichstein, M., 2008. Terrestrial ecosystem carbon dynamics and climate feedbacks. Nature 451, 289–292.

Held, I., Suarez, M., 1974. Simple albedo feedback models of the icecaps. Tellus 26(6), 613–629.

Joos, F., Meyer, R., Bruno, M., Leuenberger, M., 1999. Variability in the carbon sinks as reconstructed for the last 1000 years. Geophysical Research Letters 26(10), 1437–1440.

Lemoine, D., 2010. Paleoclimatic warming increased carbon dioxide concentrations. Journal of Geophysical Research 115(D22), D22122.

Lemoine, D. M., Traeger, C. P., 2010. Tipping points and ambiguity in the integrated assessment of climate change, department of Agricultural and Resource Economics, UC Berkeley, Working Paper Series 1704668.

Lenton, T., Held, H., Kriegler, E., Hall, J., Lucht, W., Rahmstorf, S., Schellnhuber, H., 2008. Tipping elements in the earth's climate system. Proceedings of the National Academy of Sciences 105(6), 1786.

Miranda, M., Fackler, P., 2004. Applied computational economics and finance. MIT Press Books 1.

Naevdal, E., 2006. Dynamic optimisation in the presence of threshold effects when the location of the threshold is uncertain-with an application to a possible disintegration of the western antarctic ice sheet. Dynamic optimisation in the presence of threshold effects when the location of the threshold is uncertain-with an application to a possible disintegration of the Western Antarctic Ice Sheet 30(7), 1131–1158.

Naevdal, E., 2007. The economics of the thermohaline circulation–a problem with multiple thresholds of unknown locations. Resource and energy economics 29(4), 262–283.

Nordhaus, W., 2008. A question of balance: Weighing the options on global warming policies. Yale University Press.

Rahmstorf, S., Crucifix, M., Ganopolski, A., Goosse, H., Kamenkovich, I., Knutti, R., Lohmann, G., Marsh, R., Mysak, L., Wang, Z., et al., 2005. Thermohaline circulation hysteresis: A model intercomparison. Geophys. Res. Lett 32, L23605.

Randall, D., Wood, R., Bony, S., Colman, R., Fichefet, T., Fyfe, J., Kattsov, V., Pitman, A., Shukla, J., Srinivasan, J., Stouffer, R., Sumi, A., Taylor, K., 2007. Cilmate models and their evaluation. in: Climate change 2007: The physical science basis. contribution of working group i to the fourth assessment report of the intergovernmental panel on climate change [solomon, s., d. qin, m. manning, z. chen, m. marquis, k.b. averyt, m.tignor and h.l. miller (eds.)]. Technical report, Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Rezai, A., 2010. Recast the dice and its policy recommendations. Macroeconomic Dynamics 14(S2), 275–289.

Roe, G., Baker, M., 2007. Why is climate sensitivity so unpredictable. Science 318(5850), 629–632.

Roe, G., Baker, M., 2011. Comment on "another look at climate sensitivity" by zaliapin and ghil (2010). Nonlin. Processes Geophys 18, 125–127.

Skiba, A., 1978. Optimal growth with a convex-concave production function. Econometrica 46(3), 527–539.

Tahvonen, O., Salo, S., 1996. Nonconvexities in optimal pollution accumulation. Journal of Environmental Economics and Management 31(2), 160–177.

Tahvonen, O., Withagen, C., 1996. Optimality of irreversible pollution accumulation* 1. Journal of Economic Dynamics and Control 20(9-10), 1775–1795.

Tsur, Y., Zemel, A., 1996. Accounting for global warming risks: Resource management under event uncertainty. Journal of Economic Dynamics and Control 20(6-7), 1289–1305.

Tsur, Y., Zemel, A., 1998. Pollution control in an uncertain environment. Journal of Economic Dynamics and Control 22(6), 967–975.

Tsur, Y., Zemel, A., 2009. Endogenous discounting and climate policy. Environmental and Resource Economics 44(4), 507–520.

Webb, M., Senior, C., Sexton, D., Ingram, W., Williams, K., Ringer, M., McAvaney, B., Colman, R., Soden, B., Gudgel, R., et al., 2006. On the contribution of local feedback mechanisms to the range of climate sensitivity in two gcm ensembles. Climate Dynamics 27(1), 17–38.

Zaliapin, I., 2011. Reply to roe and baker's comment on "another look at climate sensitivity" by zaliapin and ghil (2010). Nonlin. Processes Geophys 18, 129–131.

Zaliapin, I., Ghil, M., 2010. Another look at climate sensitivity. Nonlinear Processes in Geophysics 17, 113–122.

Zickfeld, K., Eby, M., Matthews, H. D., Schmittner, A., Weaver, A. J., 2011. Nonlinearity of carbon cycle feedbacks. Journal of Climate 24(16), 4255–4275.