

UNIVERSITY OF OSLO
FACULTY OF SOCIAL SCIENCES
DEPARTMENT OF ECONOMICS
PROGRAMME: MPhil ENERGY, RESOURCE AND DEVELOPMENT ECONOMICS

THE SOCIAL COST OF CARBON UNDER
UNCERTAINTY OVER CLIMATE
SENSITIVITY AND DAMAGES

MASTER THESIS
30 CREDITS
SUBMITTED BY: MARTE ELINE MENØY
JULY 2020

THE SOCIAL COST OF CARBON WHEN
LEARNING ABOUT THE CLIMATE
SENSITIVITY AND DAMAGES

OPTIMAL POLICY RESPONSE TO UNCERTAINTY OVER CLIMATE
SENSITIVITY AND CONVEXITY OF DAMAGES IN AN INTEGRATED
ASSESSMENT MODEL OF CLIMATE CHANGE.

@ Marte Eline Menøy

2020

The social cost of carbon when learning about the climate sensitivity and damages

<https://www.duo.uio.no/>

Publisher: Representralen, University of Oslo

Abstract

The current study considers how the social cost of carbon responds to joint uncertainty over climate sensitivity and damages. Climate sensitivity is the temperature response to a doubling of pre-industrial carbon concentration in the atmosphere. Uncertainty over damages is implemented over the damage exponent, and reflect that we do not know how steeply damages increase at higher temperatures. These two uncertainties interact, and therefore it is insightful to investigate their joint effect. In the literature the two have only been studied separately. The analysis is carried out in a wide-spread Integrated Assessment Model of Climate Change (IAM), more specifically a dynamic programming version of the Dynamic Integrated Climate Economy model (DICE). Uncertainty is implemented as a simplified Bayesian learning process. We simplify the learning process by using the exogenous evolution of the variance in the normal-normal Bayesian learning model. Comparing the results from the model with different uncertainties to a deterministic version of the same model, we find that both joint and single uncertainty contributes to a higher social cost of carbon (SCC). Uncertainty over climate sensitivity contributes more to the SCC in the short term, while uncertainty over the convexity of damages contributes more in the long term. The model is coded in MATLAB.

Acknowledgements

First and foremost I would like to thank my supervisor Christian Traeger for providing me with this interesting, challenging, and most of all rewarding project. His excellent guidance and encouragement has been invaluable. I will also like to give a special thanks to my co-supervisor Svenn Jensen for giving me excellent coding advise.

I am grateful to my family and friends for being supportive and motivating, and especially to my boyfriend Sølve for his love, support and interesting discussions.

Last but not least I want to thank Oslo Centre for Research on Environmentally friendly Energy (CREE) for providing me with a scholarship.

Any flaws, mistakes or inconsistencies in this thesis are entirely my own.

Contents

1	Introduction	1
2	Background	2
2.1	Integrated Assessment Models of Climate Change	2
2.2	Uncertainty in IAMs	3
2.2.1	Uncertainty over the climate sensitivity parameter	3
2.2.2	Uncertainty over the damage function	4
3	The Model	6
3.1	The exogenous processes	7
3.2	The endogenous processes	8
3.3	Uncertainty	11
3.3.1	Model implementation	11
3.3.2	Prior distributions	13
4	Solution method	15
4.1	Dynamic programming and the Bellman Equation	15
4.2	The collocation method and value function iteration.	16
4.2.1	Higher dimensional state space	18
4.2.2	Uncertainty	19
4.3	Numeric implementation	20
4.4	The Social Cost of Carbon	21
5	Results	23
5.1	Individual and joint uncertainty	23
5.2	Learning	28
5.3	Control Rules	31
6	Conclusion	35
A	Parameters	40
B	Comparison to 4 state model	41

List of Figures

1	Illustrated IAM. The oval boxes are the control variables, the rectangular boxes are the state variables, and the two red colored boxes represent the two uncertain state variables.	6
2	Comparison plots over the SCC and abatement rate. The plots compare model runs for joint uncertainty, only uncertainty over the convexity of damages, only uncertainty over the climate sensitivity and finally the deterministic case. The top figures shows the time paths simulated for the first 35 years, while the bottom figures shows the same over a 180 years time horizon	24
3	The mean, median, expected path and 67% confidence interval over the SCC simulated by 10 000 random paths 100 years into the future. The graph show this simulation for joint uncertainty and a shock variance of 20.	26
4	Optimal SCC for different prior distributions over the climate sensitivity. Compares prior variances of 2.47 and 3 to the deterministic SCC.	27
5	Comparison plots over the SCC. Compares model runs for two different speed of learning. More specifically the plots depicts the difference σ_{ξ}^2 of 2 versus 20 yields compared to the deterministic case.	29
6	Conditional variance over climate sensitivity and damage convexity. For σ_{ξ}^2 equal 2 and 20.	30
7	Control rules for the SCC over climate sensitivity (CS) and the damage convexity (b2). The blue dots marks the expected path value for the year 2040 and 2240.	31
8	Difference in control rules for joint and single uncertainty runs at year 2040 (figure a and b) and 2240 (figure c and d). Depicts optimal SCC for different values of S and b_2 . The other state variables are held fixed at their 2040 expected value for figure a and b, and at their year 2240 expected value for figure c and d.	33
9	Comparison plots over the SCC and abatement rate. Compares the deterministic 6 state model to the deterministic 4 state model. The top figures shows the time paths simulated for the first 50 years, while the bottom figures shows the same over a 200 years time horizon	41

List of Tables

1	Starting values, number of nodes and intervals over the state variables.	21
2	Numeric values for the SCC for selected years from figure 2. In parentheses is the percent increase from deterministic SCC to uncertainty SCC.	24
3	Numeric values for the abatement rate for selected years from figure 2 .	25
4	SCC for plots in figure 5, and extended for years not shown in the graph. Only uncertainty over the climate sensitivity = CS. Only uncertainty over the convexity of damages = D.	29
5	Model parameters	40

1 Introduction

The planet is getting warmer from greenhouse gas emissions, that much seems certain. At the very least, it is "extremely likely", as The Intergovernmental Panel on Climate change (IPCC) puts it. Yet, there is a lot we do not know about how climate change will affect us; both future economic and climatic consequences are governed by vast uncertainties. How to find policies that respond optimally to these uncertainties is an important and hot topic in the field of economics today.

Integrated assessment models (IAM) are frequently applied in studies of the economic consequences of climate change. This class of dynamic models incorporate how the economy and the climate interact. Thus, they make a handy tool for understanding the uncertainty in this context. In the current study, we use a widely applied IAM; the Dynamic Integrated Climate Economy model (DICE). However, we use a dynamic programming version of the model to study how learning over two key parameters affects optimal policy.

The current study considers how the social cost of carbon respond to joint uncertainty over climate sensitivity and damages. Climate sensitivity tells the temperature response to a doubling of pre-industrial carbon concentration in the atmosphere. Uncertainty over damages is implemented over the damage exponent, and reflect that we do not know how steeply damages increase at higher temperatures. These two uncertainties interact and therefore it is insightful to investigate their joint effect. In the literature the two have only been studied separately. The analysis is carried out in a wide-spread Integrated Assessment Model of Climate Change (IAM), more specifically a dynamic programming version of the Dynamic Integrated Climate Economy model (DICE). Uncertainty is implemented as a simplified Bayesian learning process. We simplify the learning process by using the exogenous evolution of the variance in the normal-normal Bayesian learning model. The variance falls exogenously over time. This means that we abstract away from possible active learning, meaning that the decision maker cannot emit more CO₂ to faster gain knowledge over the climatic system. However, Jensen and Traeger (2016) shows that this is of minor importance.

2 Background

This section starts by introducing Integrated assessment models of climate change (IAM). Then we will go on to discuss uncertainty implemented in such models. Finally, we briefly review the literature on uncertainty in dynamic programming implementations of IAMs. More specifically, we explore uncertainties governing 1) the damage function, and 2) the climate sensitivity.

2.1 Integrated Assessment Models of Climate Change

Integrated assessment models of climate change combines the scientific and socio-economic aspects of climate change to assess climate policies (Kelly & Kolstad, 1999b). The intergovernmental panel of climate change (IPCC) defines IAMs as "a method of analysis that combines results and models from the physical, biological, economics and social sciences, and the interactions among these components in a consistent framework to evaluate the status and the consequences of environmental change and the policy responses to it"(IPCC, n.d.). Integrated assessment models essentially model the dynamic interactions between the climate and the economy. Researchers aim to understand the underlying forces at play to inform policy makers. In the U.S., IAMs are used to price greenhouse gas emissions in cost benefit analyses of federal policies (Greenstone, Kopits, & Wolverton, 2013). The models are assessed on a regular basis by the IPCC.

Within the class of integrated assessment models, there are significant variations between models with respect to which assumptions they build on and the level of complexity. An example of a complex IAM is the Global Change Assessment Model (GCAM). This model represents in a detailed manner the behaviour and interactions of five different systems; the energy system, water, agriculture and land use, the economy and climate. The purpose of this rather complex model is to explore future scenarios and quantify the implications of possible future conditions (JGCRI, n.d.). However, less complex cost benefit oriented models such as The Climate Framework for Uncertainty, Negotiation and Distribution (FUND), Policy Analysis of the Greenhouse Effect (PAGE) and The Dynamic Integrated Model of Climate and the Economy (DICE), are frequently used. The current study applies a version of the latter. The benefit of the less complex IAMs is that they are more transparent, and so one can more easily understand the different drivers of policy.

The DICE model is developed by William Nordhaus. In 2018 he was awarded the Nobel

Memorial Prize in Economics Sciences "for integrating climate change into long-run macroeconomics analysis"(Nobel Media AB 2020, n.d.). DICE extends the classic Ramsey growth model to include a climate module. The first version of the model was published in 1992 (Nordhaus, 1992), and has since been updated several times, most recently in 2016 (Nordhaus, 2017). The current study applies a dynamic programming implementation of the DICE-2013 (Nordhaus, 2013)(Traeger, 2014). The model is outlined in section 3.

2.2 Uncertainty in IAMs

The first attempts to analyse the implications of uncertainty in IAMs mainly applied Monte Carlo methods. However, this approach turned out to have some shortcomings. First, Monte Carlo methods do not allow for feedback, and so the anticipation of learning and the adaption of future policies to future shocks cannot be taken into account. Second, the Monte Carlo approach has been shown to possibly get the sign of the uncertainty effect wrong (Crost & Traeger, 2013).

In recent years, recursive dynamic programming implementations of IAMs have got more attention. The approach was pioneered by Kelly and Kolstad in 1999, but by 2013 there were only two other studies conducted based on this method. Traeger (2014) and with various co-authors took the lead in this method the last decade (Crost & Traeger, 2013), (Crost & Traeger, 2014), (Jensen & Traeger, 2014) and (Traeger & Lemoine, 2014). Today, analysing uncertainty in IAMs using recursive methods is a popular approach (Lemoine & Rudik, 2017).

The remainder of this section reviews the relevant literature on uncertainty over the climate sensitivity parameter and the parameter for damage convexity. It is important to note that the models in the cited articles differ, and so they do not necessarily represent comparable experiments.

2.2.1 Uncertainty over the climate sensitivity parameter

Climate sensitivity describes the relationship between the level of CO₂ in the atmosphere and global temperature. More specifically, it is defined by how many Celsius degrees of global warming that result from a doubling of the CO₂ concentration in the atmosphere, as compared to pre-industrial times. Before the Industrial Revolution, global average atmospheric CO₂ concentration was about 280 ppm¹. In april 2020 the Mauane Loa Observatory in Hamwahi measured the CO₂ concentration to be 416.21

¹parts per million

ppm (National Oceanic and Atmospheric Administration, 2020). Reducing the uncertainty over the climate sensitivity parameter has been an important part of climate research over the years. Yet, the progress has been modest. The first IPCC report from 1990 used a likely range over climate sensitivity of 1.5°C to 4.5°C (IPCC, 1990). The same range was used in the Charney report, which is considered to be one of the first comprehensive works on the assessment of climate change due to CO₂ emissions (Charney, 1979). The most recent IPCC report still use the same range (IPCC, 2013). However, the estimate is now based on more substantial evidence and we know more about the climatic system today (Knutti, Rugenstein, & Hegerl, 2017).

As the true value of the climate sensitivity parameter is still unknown, it is useful to gain an understanding of how to meaningfully incorporate the uncertainty into the cost-benefit framework of integrated assessment models. Uncertainty over the parameter has been studied in various settings. In the following, we review some of the studies that have been conducted within a similar framework to the one at hand.

Kelly and Zhuo (2015) and Hwang, Reynès, and Tol (2017) study learning under a fat tailed uncertainty distribution. Both studies find that the anticipation of learning reduces optimal emission control compared to the case with only uncertainty and no learning. The thicker the tail, the larger counteracting learning effect.

Fitzpatrick and Kelly (2017) evaluates the common policy recommendation of a 2°C temperature target with uncertainty and learning over the climate sensitivity. Like the current study, Fitzpatrick and Kelly also use a dynamic programming version of the DICE model. Their results suggest that in the short run it is difficult to meet the 2°C target, but with learning the social planner can achieve the target over time with sustained abatement effort. One implication of their findings is that a stringent target can induce welfare losses. That is because as we learn new information about the climate system another target may be optimal.

2.2.2 Uncertainty over the damage function

Quantifying economic damages from a warmer climate is an important factor in estimating the Social Cost of Carbon. Quantifying in a meaningful way how the costs of abatement stand to the benefits of abatement is important for policy decisions today. The problem is that we don't know the economic cost of global warming, and therefore it is important to gain an understanding of how this uncertainty affects policy making. This section reviews papers that investigate this in a dynamic programming framework.

Crost and Traeger (2014) looks at damage uncertainty as stochastic shocks over the two parameters governing the levels and the convexity of damages. They study a high- and a low-level scenario, and run both using expected utility and Epstein-Zin-Weil utility. Their results indicate that uncertainty about the level of damages slightly reduces the optimal abatement rate and optimal carbon tax. The optimal abatement level is also higher under Epstein-Zin-Weil utility than under expected utility.

The perhaps most thorough examination of uncertainty over the damage function in a dynamic programming framework, was undertaken by Rudik (n.d.). The author analyses uncertainty over the DICE damage function using four different approaches. First, in what is named the "uncertainty framework", the probability distribution is implemented over the level- and the convexity-parameters. Second, in the "learning framework", Bayesian learning is applied to the convexity-parameter. Third, in the "robust control framework", he allows for the decision maker to be uncertain about both the levels and convexity of damages, while also being aware that the damage function itself may be misspecified. This implies that the social planner insures against catastrophic outcomes. Finally, in the fourth framework he combines learning over the damage exponent with robust control. The author finds that without learning the total effect of uncertainty initially decreases the carbon tax by a small amount. Over time, however, it increases the carbon tax by an amount up to 1%. When analysing uncertainty with learning, the effect on the carbon tax is always that it increases upwards, up to 5%. Robust control without learning increases the optimal carbon tax by 1-2%, and slightly decreases the tax with learning.

Robust control incorporates the fact that we do not know the underlying form of the damage function, but it also guards against unknown misspecifications, and so guard against catastrophic events. This has the potential to induce big welfare losses ex ante if it turns out that the damage function was correctly specified. Rudik (n.d.) estimate that welfare loss can be 215 billion dollars in this case. He also find that learning about the temperature elasticity of damages can produce welfare gains also in the amount of hundreds of billions, and therefore state that research aimed at updating damage functions has the potential to be valuable.

3 The Model

Our model is a dynamic programming implementation of the 2013 version of the DICE model (Nordhaus, 2013). The dynamic programming adaption is from Traeger (2014), which originally is based on the 2007 version of the DICE model Nordhaus (2007). Here the model is updated to reflect the 2013 DICE model. Traeger (2014) reduces the state space of the DICE model to only four state variables. Having as few as possible state variables is useful because of the curse of dimensionality in dynamic programming. Implementing uncertainty such as Bayesian learning in a dynamic programming model requires additional state variables. The present study uses two additional state variables to implement learning over the climate sensitivity and the convexity of damages, this will be outlined in section 3.3.

The main relationships in the model is illustrated in figure 1. World output is a function of the state variables capital, CO2 stock and temperature, as well as exogenous variables such as labour and technological growth. Time is also a state variable, as this captures time evolving exogenous processes without adding additional state variables. Capital is used as an input in production, a bi-product of production is CO2 emissions which accumulates in the atmosphere and increases the global temperature. Higher temperatures leads to economic damages that decreases production. Output is either invested in more capital or consumed. From consumption we get social welfare. Without the emissions-temperature-economy interaction this model represent a standard Ramsey growth model.

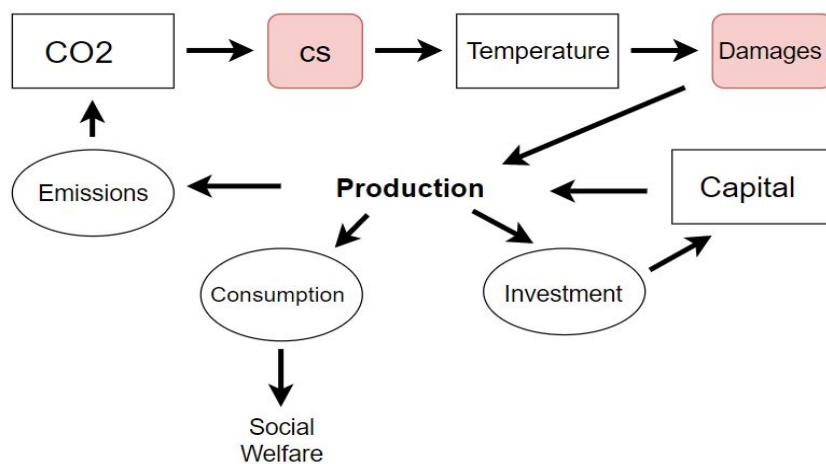


Figure 1: Illustrated IAM. The oval boxes are the control variables, the rectangular boxes are the state variables, and the two red colored boxes represent the two uncertain state variables.

3.1 The exogenous processes

The population/labor in each period is defined as:

$$L_t = L_0 \frac{L_\infty}{L_0}^{-(1-g_{L^*})^t} \quad (1)$$

where L_0 is the initial population, L_∞ is the asymptotic population and g_{L^*} is the rate of convergence to asymptotic population.

The technology level in each period is defined as:

$$A_t = A_0 \exp \left[g_{A,0} \frac{1 - \exp(-\delta_A t)}{\delta_A} \right] \quad (2)$$

where A_0 is the initial technological level in the economy, $g_{A,0}$ is the initial growth rate of technology and δ_A is the growth rate's constant rate of decline.

The carbon intensity of production is assumed to decrease exogenously, and is defined as:

$$\sigma_t = \sigma_0 \exp \left[g_{\sigma,0} \frac{(1 + \delta_\sigma)^{t-1}}{\ln(1 - \delta_\sigma)} \right] \quad (3)$$

where σ_0 is the initial level of the carbon intensity in production, $g_{\sigma,0}$ is the initial negative growth rate, and δ_σ is the growth rate's rate of decrease.

The abatement cost coefficient is an exogenous function of backstop technology.

$$\Psi_t = \left[\frac{a_0 \sigma_t}{a_2} \right] \exp(-g_\Psi t) \quad (4)$$

Abatement of emissions costs, but due to technological progress it becomes cheaper as time goes by, this is represented by a_0 which is the initial cost of backstop technology (in 2013). a_2 is the cost exponent and g_Ψ is the growth rate of abatement cost. As time goes by the expression $\exp(-g_\Psi t)$ goes towards zero, which means that we have reached the backstop technology, where energy from renewable energy has become cheaper than non-renewables, and so makes the latter obsolete.

CO2 emissions from land use change and forestry (LUCF)² also follows an exogenous

²Emissions from human activities which either change the way land is used or has an effect on the biomass in existing biomass stock (Pitesky, Stackhouse, & Mitloehner, 2009).

process and is defined as:

$$B_t = B_0 \exp(-\delta_B t) \quad (5)$$

CO2 is the most important of the greenhouse gases but other types of greenhouse gases also contributes to radiative forcing:

$$EF_t = EF_0 + 0.01(EF_{100} - EF_0) \times 10 + \min(t, 90) \quad (6)$$

where EF_0 is external forcing in year 2000 and EF_{100} is external forcing in year 2100.

In DICE the carbon cycle is divided into three levels. The first level is the atmosphere which absorbs emissions, but as the carbon concentration increase some carbon goes into the second level, which is the lower biosphere and the shallow ocean. Third, carbon from the second level escapes into the deep ocean. When emissions increases, these boxes fill up. When the ocean reaches the limit of carbon uptake, we loose these natural carbon sinks and emissions go straight into the atmosphere. In this model these three dynamics are replaced by an exogenous process:

$$\delta_{M,t} = \sigma_0 \exp\left[\delta_{M,0} \frac{(1 + \delta_\sigma)^t - 1}{\ln(1 + \delta_\sigma)}\right]. \quad (7)$$

3.2 The endogenous processes

The model contains six state variables, in this section the state variables capital K_t , atmospheric carbon stock M_t , temperature T_t and time t_t will be explained. In section 3.3, we will discuss the two additional state variables, the uncertainty process over the climate sensitivity and the damage exponent.

The model uses a Cobb-Douglas production function, that makes gross output a function of labour L_t , technology A_t and capital K_t : $Y_t^{gross} = (A_t L_t)^{1-\kappa} K_t^\kappa$. As K_t changes a lot over time, approximating it becomes computationally expensive. That is why the model normalizes capital to effective labour units, $k_t = \frac{K_t}{A_t L_t}$. Which gives per capita gross production, $y_t^{gross} = \frac{Y_t^{gross}}{A_t L_t} = k_t^\kappa$. The same argument for capital goes for consumption as well, which is also expressed in effective labor units, $c_t = \frac{C_t}{A_t L_t}$.

The interaction between the economy and the climate system is defined through net production which is defined as:

$$y_t^{net} = \frac{1 - \Lambda(\mu_t)}{1 + D(T_t)} k_t^\kappa = \frac{1 - \Psi_t \mu_t^{a2}}{1 + b_1 T_t^{b2}} k_t^\kappa \quad (8)$$

Where y_t^{net} is decreasing in abatement cost, $\Lambda(\mu_t)$, and damages from increased temperature, $D(T_t)$.

The abatement cost is defined as:

$$\Lambda(\mu_t) = \Psi_t \mu_t^{a_2} \quad (9)$$

where μ_t is the emission control rate, or abatement rate, $\mu_t \in [0, 1]$, and it characterizes the emissions avoided under a climate policy compared to a "business as usual" scenario. a_2 is the cost exponent and Ψ_t is the abatement cost coefficient, and is defined in section 3.1 in equation (4).

The damage function is defined as a fraction of gross output:

$$D(T_t) = b_1 T_t^{b_2} \quad (10)$$

where b_1 is the level of economic damages due to temperature increase over temperatures in 1900, and b_2 is the convexity of damages. In this thesis the mean over b_2 is uncertain and a state variable, how this is implemented is explained in section 3.3.

Time is also a state variable in this model. Having time as a state variable makes it possible to approximate the value function with the time-varying exogenous parameters without making them state variables. Time is natural unbounded, and therefore the model uses a monotonic transformation, mapping time on the interval $[0, 1)$ in equation (11), (Traeger, 2014).

$$\tau = 1 - \exp[-\zeta t] \quad (11)$$

τ is referred to as artificial time, while t is real time. As will be explained more in-depth in section 4 the model's objective function is approximated at Chebychev nodes, and the numerical parameter ζ decides how close to the early real time periods the nodes concentrate. This thesis uses a value of 0.02, which is the same as in Traeger (2014). The smaller the value of ζ , the closer to the early time periods the nodes will be centered. It is useful to have it set to a small value as the DICE model states that most action will be taken early in the time horizon (Traeger, 2014).

The equation of motion for capital is:

$$k_{t+1} = [(1 - \delta_k)k_t + y_t - c_t] \exp[-(g_{A,t} + g_{L,t})]. \quad (12)$$

Next period capital is the previous period's net production that is not consumed. Capital also depreciates by the annual rate of δ_k . $\exp[-(g_{A,t} + g_{L,t})\Delta t]$ reflects that the effective labour unit have an annual growth rate of $g_{A,t} + g_{L,t}$.

Anthropogenic emissions is defined as:

$$E_t = (1 - \mu_t)\sigma_t A_t L_t k_t^\kappa + B_t \quad (13)$$

and is the sum of industrial emissions (first part) and emissions from land use change and forestry, B_t . σ_t is the emission intensity of production.

We assume today's emissions add directly to the next period's atmospheric carbon stock, which is defined as:

$$M_{t+1} = M_{pre} + (M_t - M_{pre})(1 - \delta_{M,t}) + E_t \quad (14)$$

M_{pre} is the pre-industrial level of CO2 in the atmosphere, and so is the steady state level in the absence of anthropogenic emissions. The carbon stock in the next period is the sum of the pre industrial steady state level of carbon in the atmosphere and the excess amount of carbon in the previous period. $\delta_{M,t}$ is the rate of carbon removal from the atmosphere, represented in equation (7). Not all CO2 that is emitted goes straight into the atmosphere, some also go into the shallow ocean and the biosphere. The rate of this process is not stationary, its time and state dependent, as when the excess amount of carbon grow the carbon reservoirs get filled up and so the ability to take up more carbon decreases. Therefore $\delta_{M,t}$ is decreasing in time and the carbon stock.

Global average temperature change is a delayed response to forcing from the atmospheric carbon stock M_t , and other non CO2 sources, EF_t . The equation of motion for temperature is defined as:

$$T_{t+1} = (1 - \sigma_{forc})T_t + \sigma_{forc}S_t \left[\frac{\ln \frac{M_t}{M_{pre}}}{\ln 2} + \frac{EF_t}{\eta_{forc}} \right] - \sigma_{ocean}. \quad (15)$$

Radiative forcing is the difference between the energy going into the earth's atmosphere and the energy going out. This energy imbalance leads to warming. Before industrial times radiative forcing was quite stable, but an increased concentration of greenhouse gases in the atmosphere stops more energy from going back out from the earth's atmosphere, and so the radiative forcing increases. Radiative forcing is logarithmic in the atmospheric carbon stock M_t and EF_t is radiative forcing from other

sources than CO₂. Climate sensitivity, S_t represent the equilibrium temperature response to a doubling of CO₂ in the atmosphere compared to pre-industrial levels. The parameter also captures that we assume warming to be approximately proportional to radiative forcing, and that the short term temperature response is proportional to climate sensitivity. The parameter σ_{forc} captures that warming happens with a delay. Another process is represented in σ_{ocean} , and shows the cooling effect the ocean has on global temperature. The ocean heats up slower than the atmosphere, and while the ocean is catching up, it also cools down/slow the heating process.

3.3 Uncertainty

This section explains our implementation of the uncertainty governing climate sensitivity mean, S_t , and the mean over the damage exponent $b_{2,t}$. First, we discuss how we imitate a Bayesian normal-normal learning process over the parameters, then we go on to discuss the chosen prior probability distributions for each variable.

3.3.1 Model implementation

We use a modified first-order auto-regressive process, AR(1), to imitate a Bayesian normal-normal learning process. In a normal-normal learning process both the prior, governing the decision maker's subjective uncertainty, and the shock, representing nature's stochasticity, are normally distributed. As a result, also the posterior after Bayesian updating remains normally distributed. Thus, the evolution of uncertainty can be captured by two parameters, the mean and the variance of the subjective distribution. A convenient feature of the standard normal-normal Bayesian learning model is that the variance falls over time independently of the signal realizations. Thus, we can model it as an exogenous dependent process and avoid the use of an additional state variable. Kelly and Kolstad (1999a) show for the case of climate sensitivity that a more comprehensive updating would result in a normal-normal Bayesian learning model where the evolution of the variance depends not only on time but also on the evolution of the carbon stock. In principle, the decision maker can speed up learning by emitting more carbon (active learning). However, Jensen and Traeger (2016) show that the ability to actively increase the learning speed by increasing GHG emissions will not play a quantitatively meaningful role in climate change. Thus, here we neglect such active learning, which saves two state variables in an already high dimensional dynamic programming problem.

We introduce a state variable for the mean of the climate sensitivity and the mean damage exponent. Since both states follow a similar process, we let Υ_t represent either

the mean of S_t or $b_{2,t}$. The construction ensures that the long-term distribution over the parameters converges to the decision maker's prior (subjective) uncertainty. In detail, the state's equation of motion resembles a modified weighted AR(1) process:

$$\Upsilon_t = \gamma_t \Upsilon_{t-1} + (1 - \gamma_t) \tilde{\omega}_t \quad \text{where } \Upsilon_t \in \{S_t, b_{2,t}\} \quad (16)$$

where both the auto-regressive coefficient and the distribution over $\tilde{\omega}_t$ evolve over time to match the mean's evolution under the normal-normal Bayesian learning model.

In a normal-normal Bayesian learning model the variance over the subjective prior falls exogenously over time. This process is characterized by the precision, ν_t , which grows linearly in equation (17). ν_t is the inverse of the subjective prior variance at time t , σ_t^2 .

$$\nu_{t+1} = \nu_0 + t\nu_\varepsilon. \quad (17)$$

The precision is a combination of two variances. Where $\nu_0 = \frac{1}{\sigma_0^2}$, represent the precision of our prior in period 0, which is our initial belief of the parameter value. $\nu_\varepsilon = \frac{1}{\sigma_\varepsilon^2}$, characterizes the precision over the noise that slows learning³.

The evolution over the mean of the state variables are governed by the realized stochastic shock in each period $\tilde{\omega}_t$ and the AR(1) factor γ_t . The predictive distribution over $\tilde{\omega}_t$, has a variance that is dependent on the variance over the prior distribution in period t and the variance over the stochastic shock.

$$\sigma_{\omega,t}^2 = \sigma_\varepsilon^2 + \frac{1}{\nu_t}. \quad (18)$$

The AR(1) factor is defined in terms of the precision (17), and evolves over time in the following manner:

$$\gamma_t = \frac{\nu_t}{\nu_t + \nu_\varepsilon}. \quad (19)$$

Equation (19) and (18) govern the learning process in equation (16). The AR(1) factor increases asymptotically towards unity and the predictive variance decreases asymptotically towards the shock variance. Both are calculated so that standing in

³The current model is a simplified and also a more abstract learning model. If we literally had an Bayesian learning model, we would learn about the parameters by observing the climate and the economy. Hence the variance would in that case depend on the evolution of the climate and the economy.

period 0 the decision maker believe the real value of the parameters to lie within the prior distribution in period 0, σ_0^2 . Moving forward in time ν_t ensure that learning takes place. This exogenously decided learning process let's us decide how fast the decision maker learns. The parameter that let us do that is the variance over the stochastic noise, σ_ε^2 . This will be discussed further in section 5.2.

3.3.2 Prior distributions

The DICE damage function is stated in equation (10),

$$D(T_t) = b_1 T_t^{b_2}$$

And from equation (8) we have that damages can be expressed as percentage damages of gross product given a temperature. The percentage of damages is then equal to $\frac{k^\kappa}{1+D(T_t)}$. The DICE model assumes a quadratic form, so that $b_2 = 2$. b_1 is estimated by fitting the quadratic function to a set of existing damage estimates.

A probability distribution is implemented only over the damage exponent. That is because this is considered the more interesting parameter as it changes the convexity of the damage function, and so we expect it to be the most relevant for the uncertainty effect. We use a prior distribution over b_2 taken from Rudik (n.d.). To estimate the distributions over the damage parameters he uses the regression in equation (20), which is based on data from a recent metastudy by Howard and Sterner (2017). In this study

$$\ln D(T_i) = \ln b_1 + b_2 \ln T_i + \epsilon_i \quad (20)$$

The estimated \hat{b}_2 have a normal distribution with a mean = 1.88 and a variance = 0.203. Rudik (n.d.) uses this as the prior over the damage exponent in his implementation of Bayesian learning, this thesis will do the same.

The prior distribution over the climate sensitivity is constructed by fitting a normal distribution to reflect a 66 % confidence interval of 1.5°C - 4.5°C, with a mean value of 3°C. Which reflects what the latest published IPCC report (IPCC, 2013) uses as the likely interval over the climate sensitivity. As discussed in chapter 2 the likely range over the climate sensitivity has not changed much over the years. There's many studies that try to estimate the climate sensitivity and make the likely interval over the parameter smaller. Knutti et al. (2017) looks into different studies that aim to estimate the probability distribution over the parameter. They find that their overall assessment is broadly consistent with the IPCC, and that the likely value of the climate

sensitivity is 3°C.

Based on this the chosen prior over climate sensitivity is a mean equal to 3°C and a variance of 2.47°C.

The variance that governs the noise in the model, σ_ε^2 , has a value of 20 for both variables. σ_ε^2 slows learning, a lower value means less noise and hence, faster learning. To see how different speed of learning affects optimal policy, a value of 2 will also be employed to represent fast learning.

4 Solution method

This thesis solves a recursive dynamic programming implementation of the DICE model. More specifically we solve the Bellman equation using the collocation method and value function iteration. This segment will explain the solution approach in detail.

4.1 Dynamic programming and the Bellman Equation

The intuition behind dynamic programming can be summarized in the following excerpt from Bellman (1957, p. 83).

"An optimal policy has the property that whatever the initial state and initial decisions are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision."

The method breaks a problem into simpler sub-problems in a recursive manner, in our problem it breaks it down into a trade-off between welfare from consumption now and from consumption in the next period. A dynamic programming system consists of state variables and control variables. The state variables determine the economic system in each period, and the decision maker will in each period optimize policy by choosing the level of the control variables.

The benefit of this approach in solving dynamic models under uncertainty is that it creates what we can call a closed loop system. In each period the decision maker will be aware of the system's state and make decision based on that. For our case with learning, the decision maker will in each period be aware of which state the system ended up in, what the remaining uncertainty is, and use that as a basis for decision making.

The dynamic programming problem is defined in equation (21). It is constructed in the form of a Bellman equation, and we assume standard preferences. The right hand side of equation (21) is the value function $V(\Theta_t)$ ⁴. The left hand side is the instantaneous time t welfare, represented by the utility function in period t , and the discounted future welfare. We maximize over consumption c_t and emission control rate μ_t . The chosen value of the control rates decides what state the economic system will end up in today and in the following period.

⁴Where Θ_t represents all six state variables in equation (21)

$$V(k_t, M_t, t_t, T_t, \tilde{S}_t, \tilde{b}_{2,t}) = \max_{\mu_t, c_t} \frac{c_t^{1-\eta}}{1-\eta} + \beta_t \mathbf{E}[V(k_{t+1}, M_{t+1}, t+1, T_{t+1}, \tilde{S}_{t+1}, \tilde{b}_{2,t+1})] \quad (21)$$

The state variables are per capita capital, k_t , atmospheric carbon stock, M_t , temperature, T_t , time, t_t , the damage exponent, b_2 , and the climate sensitivity, S_t . The decision maker chooses optimal policy by choosing the optimal level of the control variables, which is the emission control rate, μ_t , and per capita consumption c_t . The utility discount factor is β_t and incorporates the time changing population and technological growth.

$$\beta_t = \exp(-\delta_u + g_{A,t}(1-\eta) + g_{L,t}). \quad (22)$$

The time dependence is a consequence of the time invariant growth rates, and does not indicate a time inconsistent objective function. The factor β_t determines the contraction of the Bellman equation.

4.2 The collocation method and value function iteration.

We start by representing the Bellman equation (21) by a simplified version:

$$V(\Theta) = \max_x [U(x) + \beta V(g(\Theta, x))] \quad (23)$$

where Θ represent the state variables, and x is the control variables. The function $g(\Theta, x)$ is the state transition function, which simply state that the next period's states are a function of the prior period's states and chosen values of control variables x . In the following Θ will represent a one dimensional state space and x represent a single control variable, this simplifies the explanation. The following explanation relies on the work by Miranda and Fackler (2002). In section 4.2.1 we will go on explaining this solution method applied on a higher dimensional state space, and in section 4.2.2 outline how uncertainty is incorporated in the model.

The approximate solution to the Bellmann equation is found by using the collocation method. First the value function is written as a combination of chosen basis functions $\phi_1, \phi_2, \dots, \phi_n$, and coefficients c_1, c_2, \dots, c_n . It is the coefficients we will solve for. Based

on this the value function can be written as the following approximation:

$$V(\Theta) \approx \sum_{j=1}^n c_j \phi_j(\Theta) \quad (24)$$

The chosen basis functions are Chebychev polynomials with corresponding Chebychev nodes and coefficients.

After choosing the basis functions, the coefficients are estimated by requiring the approximant to satisfy the Bellman equation at n collocation nodes, $\Theta_1, \Theta_2, \dots, \Theta_n$, within a chosen interval, $[\Theta_{min}, \Theta_{max}]$. The system then consists of n non-linear equations with n unknowns. The simplified Bellman equation (23) can now be written as:

$$\sum_{j=1}^n c_j \phi_j(\Theta_i) = \max_x [U(x) + \beta \sum_{j=1}^n c_j \phi_j(g(\Theta_i, x))] \quad (25)$$

where the value function is replaced by its' approximant. The value function's approximant can be expressed in vector form as the collocation equation:

$$\Phi c = v. \quad (26)$$

Where Φ is called the collocation matrix, and is a n by j matrix whose typical nj^{th} element is the j^{th} basis function evaluated at the n^{th} collocation node, $\Phi_{nj} = \phi_j(\Theta_n)$, (Miranda & Fackler, 2002). The right hand side of equation (26) is a n by one vector, and has the typical i^{th} element of equation (25) evaluated at the i^{th} state node. c is a n by one matrix, containing the estimated coefficients that satisfy equation (26). v is a column vector where the n^{th} element is the right hand side of equation (25) evaluated at each node, Θ_n . Written out equation (26) looks like this:

$$\begin{pmatrix} \phi_0(\Theta_0) & \phi_1(\Theta_0) & \dots & \phi_{j-2}(\Theta_0) & \phi_{j-1}(\Theta_0) \\ \phi_0(\Theta_1) & \phi_1(\Theta_1) & \dots & \phi_{j-2}(\Theta_1) & \phi_{j-1}(\Theta_1) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \phi_0(\Theta_{n-2}) & \phi_1(\Theta_{n-2}) & \dots & \phi_{j-2}(\Theta_{n-2}) & \phi_{j-1}(\Theta_{n-2}) \\ \phi_0(\Theta_{n-1}) & \phi_1(\Theta_{n-1}) & \dots & \phi_{j-2}(\Theta_{n-1}) & \phi_{j-1}(\Theta_{n-1}) \end{pmatrix} \begin{pmatrix} c_0 \\ c_1 \\ \vdots \\ c_{j-2} \\ c_{j-1} \end{pmatrix} = \begin{pmatrix} \hat{v}(\Theta_0) \\ \hat{v}(\Theta_1) \\ \vdots \\ \hat{v}(\Theta_{n-2}) \\ \hat{v}(\Theta_{n-1}) \end{pmatrix} \quad (27)$$

The approximation scheme then continues by the application of the value function iteration algorithm. We use either of the following iterative updating rules. If we choose $n = j$ so that the number of basis function is equal to the number of collocation

nodes we have the updating rule in equation (28).

$$c = \Phi^{-1}v. \quad (28)$$

If $n \neq j$ we have the following updating rule of equation (29), which is the familiar Ordinary Least Squares estimator.

$$c = (\Phi' \Phi)^{-1} \Phi' v \quad (29)$$

Both equation (28) and (29) is dependent on the assumption that the basis functions are linearly independent so that the inverse of the collocation matrix exists.

The function iteration proceeds as follows. First we start with a guess over the collocation coefficients, c_0 , and solve the right hand side of equation (25) for every n^{th} node. Then we fit a new value function approximation using the guess from the last iteration and so find new values for c_1 . This proceeds until the difference between c_{k-1} and c_k is small enough to satisfy a given tolerance criterion, in this problem that is set to 10^{-1} . This algorithm is called value function iteration. The theoretical foundation for this solution method is the contraction properties of the Bellman equation. The contraction mapping theorem states that when iteratively updating the value function starting from any initial guess, the sequence will eventually converge to the one unique solution. As it will converge to the unique solution asymptotically, one set a tolerance criterion as mentioned above (Judd, 1998).

4.2.1 Higher dimensional state space

This thesis uses a dimensional state space consisting of six state variables. The solution method is very similar to the one explained above when applied on a higher dimensional state space. We still use the basis functions and collocations nodes as discussed above. The difference is that we use the basis function and collocation nodes for each of the state variable as a basis to form a multidimensional tensor product. The right hand side of equation (25) now has to be solved for every point on a multidimensional grid. This is where the "curse of dimensionality" in dynamic programming occur, as the number of nodes we need to estimate the function on increases exponentially in the number of states. If we assume that the order of approximation (order of basis functions) is equal to the number of nodes in each dimension, and that we use the six states variables used in this thesis, equation (28) can now be written as:

$$c = [\Phi_k^{-1} \otimes \Phi_M^{-1} \otimes \Phi_t^{-1} \otimes \Phi_T^{-1} \otimes \Phi_S^{-1} \otimes \Phi_{b_2}^{-1}]v. \quad (30)$$

Which is the Kronecker product of the collocation matrices in each dimension.

4.2.2 Uncertainty

Under uncertainty each stage of the dynamic problem is governed by an uncertain process. We rewrite equation (23) to equation (31):

$$V(\Theta) = \max_x [U(x) + \beta \mathbf{E}V(g(\Theta, x, \tilde{\epsilon}))]. \quad (31)$$

Where $\tilde{\epsilon}$ represent a state variable governed by a stochastic process. Now each stage is decided not only by the chosen control variables, but also by an uncertain process incorporated in $\tilde{\epsilon}$. It follows that equation (25) can be written as:

$$\sum_{j=1}^n c_j \phi_j(\Theta_i) = \max_x [U(x) + \beta \mathbf{E} \sum_{j=1}^n c_j \phi_j(g(\Theta_i, x, \tilde{\epsilon}))]. \quad (32)$$

We follow the same solution procedure as above, but now we have to take the expectation at each stage of the problem. To evaluate the optimal policy response we now have to approximate the expectation integral. We do so by the use of Gauss-Legendre quadrature. As it is infeasible to evaluate the value function at all possible realizations of the random variable, Gauss-Legendre quadrature helps us to simplify the random variable $\tilde{\epsilon}$ to a discrete random variable with mass points e_l and probabilities w_l .

Bringing the explanation over to the problem presented in this thesis, uncertainty is implemented by having the shock in equation (16) being estimated by Gauss-Legendre quadrature. More specifically is the mean of climate sensitivity and the damage exponent evaluated at mass points e_l .

$$\Upsilon_{l,t} = \gamma_t \Upsilon_{l,t-1} + (1 - \gamma_t) e_l \quad \text{where } \Upsilon_t \in \{S_t, b_{2,t}\} \quad \text{and } e_l \in \{e_{l,S_t}, e_{l,b_{2,t}}\} \quad (33)$$

The expectation in equation (21) is replaced by a set of weights, w_l . We illustrate the implementation of Gauss-Legendre quadrature in this problem in equation (34).

$$V(k_t, M_t, t_t, T_t, \tilde{S}_t, \tilde{b}_{2,t}) = \max_{\mu_t, c_t} \frac{c_t^{1-\eta}}{1-\eta} + \beta_t \sum_{l=1}^L w_l [V(k_{t+1}, M_{t+1}, t+1, T_{t+1}, S_{l,t+1}, b_{2,l,t+1})]. \quad (34)$$

Here the expectation is replaced by the weights, w_l ⁵ and the state variables is now represented by $S_{l,t+1}$ and $b_{2,l,t+1}$, where the subscript l represent that we evaluate the value function on l different points within the state space interval. When generating the Gauss-Legendre quadrature nodes and weights, we assume the covariance between the two parameters to be zero.

4.3 Numeric implementation

For the numeric implementation it is more efficient to maximize over the abatement cost Λ instead of the emission control rate, μ_t (Traeger, 2014). The constraints imposed on consumption and abatement cost is:

$$\begin{pmatrix} 1 & \frac{k_t^c}{1+b_1 T_t^{b_2}} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} c \\ \Lambda \end{pmatrix} \leq \begin{pmatrix} \frac{k_t^c}{1+b_1 T_t^{b_2}} \\ \Psi_t \end{pmatrix} \quad (35)$$

where $c, \Lambda \geq 0$. Said in words this state that the sum of consumption and abatement must be smaller or equal to net output, and that abatement cost Λ must be smaller or equal to the abatement cost coefficient, ψ .

The numeric implementation is carried out in MATLAB and take use of the CompEcon toolbox from Miranda and Fackler (2002), and KNITRO for the optimization routine. The model use Chebychev polynomials with Chebychev nodes and coefficients at a total number of 25 600. The node number for each state variable is given in table 1, together with the starting values and intervals over the state variables.

⁵Here l represent weights for both variables. If we evaluate the state variables on 3 different nodes each, then w_l is a column vector of dimensions 9×1 .

State	Value	Definition
K_0	135	In trillions USD, initial global capital stock
M_0	818.985	In GtC, stock of atmospheric CO2
T_0	0.83	Temperature in °C
S_0	3	Expected value climate sensitivity
$b_{2,0}$	1.88	Expected value damage exponent
	Interval	Nodes
K	[0.5 4]	10
M	[550 2000]	5
t	[0 1]	8
T	[0 6]	4
S	[1.5 5]	4
b_2	[1.25 2.75]	4

Table 1: Starting values, number of nodes and intervals over the state variables.

The intervals presented in table 1 represent the state space over which we estimate the state variables. The interval over climate sensitivity is 1.5°C to 5°C. As discussed in section 3.3.2 the likely interval over climate sensitivity is 1.5°C to 4.5°C. The approximation interval employed in the current study reflects that, but also extend it a half degree upwards⁶ This interval is narrow, and a valuable extension of the model would be to increase this interval. The reason for why it is not, is due to convergence issues, this will be explained more in section 6. The same problematic goes for the convexity of damages as well, where a interval of 1.25°C to 2.75°C is implemented.

4.4 The Social Cost of Carbon

To abate or not to abate, or rather how much, is the question the Social cost of carbon (SCC) aims to answer. On the one hand we experience economic damages from emitting CO2 into the atmosphere and the temperature increase that leads to, while on the other hand abatement comes with a cost. To find the optimal amount of abatement we often rely on the SCC, which in this thesis is defined in equation 36:

$$SCC_t = -1000 \frac{\partial_{M_t} V}{\partial_{K_t} V} A_t L_t. \quad (36)$$

The SCC is the marginal rate of substitution between atmospheric carbon concen-

⁶In the current study a normal distribution is employed over the parameters. However, for climate sensitivity and damages a right-skewed distribution is more likely (Wagner & Weitzman, 2015).

tration and capital⁷. It tells the value of emitting CO2 that accumulates in the atmosphere, expressed in terms of capital. More CO2 in the atmosphere leads to an increase in temperature, which again leads to more economic damages, which again decreases the capital stock. The SCC represent a trade-off between abatement cost and economic damages.

The Carbon tax is often used interchangeably with the SCC. The difference between the two is that the what is most often referred to as the carbon tax is the Pigovian tax on carbon, i.e. the monetized external effect from CO2 emissions. When abatement rate is smaller than one, $\mu_t \leq 1$, the carbon tax and the SCC coincide. While when abatement rate has reached its limit, when emissions are zero, then the carbon tax only reflects the price of keeping it that way, and in this scenario the SCC may be much higher than the carbon tax.

⁷Since capital is measured in trillions of dollars and the carbon stock in billions of tons, we have to multiply the expression by 1000. The SCC is then expressed in terms of dollars per ton of carbon. $A_t L_t$ come into play because of normalization.

5 Results

The discussion is divided into three parts. First we will discuss individual and joint uncertainty, second learning and finally the control rules.

5.1 Individual and joint uncertainty

This section presents time paths where we assume that nature draws the expected value in every period, but the social planner acts as if the world was uncertain. The difference between the deterministic path and the uncertainty paths reflects the social planner's awareness of uncertainty. More specifically, we construct the time paths by simulating the optimized value function (34) forward in time assuming the uncertain state variables to be at their expected value.

The result presented in figure 2 compares the policy implications of joint to single uncertainty. Joint uncertainty means that we have uncertainty over both parameters, while we have single uncertainty when uncertainty is only implemented over each of the variables separately.

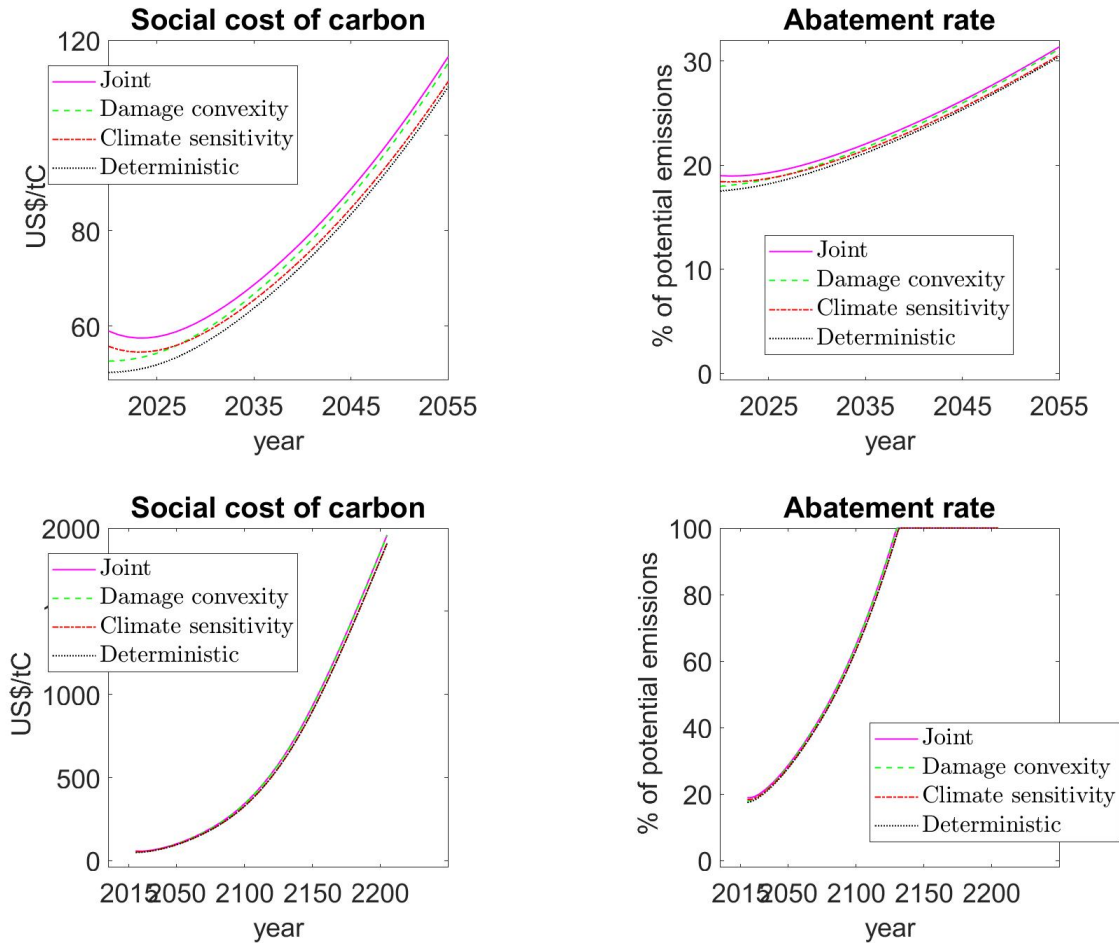


Figure 2: Comparison plots over the SCC and abatement rate. The plots compare model runs for joint uncertainty, only uncertainty over the convexity of damages, only uncertainty over the climate sensitivity and finally the deterministic case. The top figures shows the time paths simulated for the first 35 years, while the bottom figures shows the same over a 180 years time horizon

Year	Deterministic	Joint	Climate sensitivity	Damage convexity
2020	50\$	59\$ (17.3%)	56\$ (11%)	53\$ (4.7%)
2025	51\$	57\$ (12.1%)	55\$ (6.4%)	54\$ (4.7%)
2030	55\$	61\$ (9.3%)	58\$ (4%)	58\$ (4.7%)
2070	159\$	167\$ (5.3%)	160\$ (0.7%)	166\$ (4.5%)
2120	493\$	515\$ (4.5%)	494\$ (0.3%)	513\$ (4%)
2220	2188\$	2246\$ (2.6%)	2192\$ (0.2%)	2243\$ (2.5%)

Table 2: Numeric values for the SCC for selected years from figure 2. In parentheses is the percent increase from deterministic SCC to uncertainty SCC.

Year	Deterministic	Joint	Climate sensitivity	Damage convexity
2020	17.5%	19%	18.4%	18%
2025	18%	19.1%	18.6%	18.5%
2030	19.2%	20.1%	19.6%	19.7%
2070	38.8%	40%	38.9%	39.7%
2120	83.7%	85.8%	83.9%	85.6%
2220	100%	100%	100%	100%

Table 3: Numeric values for the abatement rate for selected years from figure 2

It is not easy to separate the graphs in the plots over the long time horizon in figure 2, therefore table 2 and 3 present the numeric values for these graphs.

Uncertainty over climate sensitivity and the damage convexity contributes to a higher SCC in a given period if equation 36 is a convex function of climate sensitivity and damage convexity in that period. More specifically if $\mathbf{E}SCC(b_{2,t}, S_t) > SCC(\mathbf{E}(b_{2,t}, S_t))$. This is what is known as Jensen’s inequality. From figure 2 we can see that both variables separately and together makes the case for a higher SCC compared to the deterministic SCC.

When comparing the two types of uncertainty in figure 2, we notice that the climate sensitivity alone results in a higher SCC and abatement rate at the beginning of the period than the convexity of damages does on its own. This however, changes in year 2029. One hypothesis that can explain the initially small effect from damage uncertainty is the fact that in the beginning of the period we only experience low temperatures, close to one degree. At one degree Celsius the parameter $b_{2,t}$ does not have any effect. Standing in period 1 it is only the effect from the expected impact over the course of the future that affect the SCC. In the future we will experience higher temperatures, and so the uncertainty over damages becomes more important. This is clear in table 2 where in 2120 only damage uncertainty SCC is 4.1% higher than the deterministic SCC, while only climate sensitivity uncertainty is only 0.3% higher. We will revisit this hypothesis in section 5.2.

A further study of figure 2 and table 2, show that the difference in the SCC decreases over time. This reflects the learning process. We use the value function optimized under uncertainty and simulated forward in time assuming nature draws the parameter’s

expected value in all periods. We implement an exogenous declining variance over both state variables, and so the value function is evaluated under a very low uncertainty at the high time nodes. The difference between the deterministic path and the uncertainty paths reflects the decision maker's reaction to uncertainty. Over time the variance fall and so reflect that we learn. The convergence towards the deterministic path reflects that⁸.

The abatement rate and SCC increases over time, which is similar to the structure of the DICE model. In the DICE model the optimal emissions reduction follows a "policy ramp" where we start out with a modest abatement effort and then over time increase the reductions of our GHG emissions (Nordhaus, 2007).

So far we have only depicted the expected paths. Figure 3 shows how the expected path relates to other possible time paths under uncertainty. This is done by randomly simulating our model forward in time 10 000 times. The starting values for S_t and $b_{2,t}$ is their expected value in year 2020. Figure 3 show the resulting mean, median and the 67% bound for the SCC.

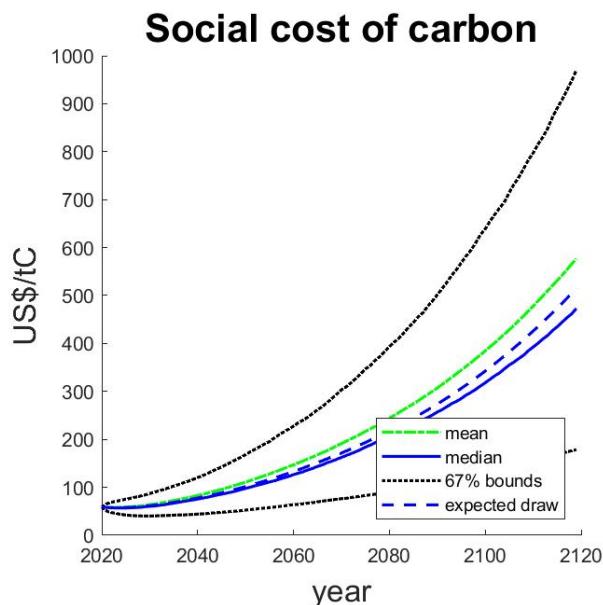


Figure 3: The mean, median, expected path and 67% confidence interval over the SCC simulated by 10 000 random paths 100 years into the future. The graph show this simulation for joint uncertainty and a shock variance of 20.

⁸The paths might not converge perfectly because by the time uncertainty resolves, the uncertainty paths have a past of different decisions leading to different states in the present.

In figure 3 we can see that the optimal SCC for 67% of the random paths lies between approximately 200\$ and 1000\$. Which means that seen from today's perspective there is a 67% chance that the optimal SCC will lie between 200\$ and 1000\$ in year 2120. This reflects that from today's perspective the climatic system is very uncertain. Comparing this result to the expected paths in figure 2 it may seem surprising that the latter does not reflect more effect from uncertainty. This can be explained by how fast learning is resolved within this model. Figure 3 reflects the uncertainty the decision maker faces in the present, while figure 2 reflects how the planner responds to uncertainty at each point in time. The variance over both parameters fall quickly. This means that while the decision maker faces much uncertainty now, he will be much more certain about the true value of the parameters in the near future.

Studying figure 3 we also notice that the mean of the random paths are over the expected time path, and the median of the random simulation lies beneath it. This can be explained by the SCC being more convex or responsive to high realizations of the parameters compared to lower realizations.

In figure 4 the effect two different prior variances have on optimal policy. The effect is only shown for only uncertainty over the climate sensitivity. Here we notice that a higher prior contributes to a higher SCC.

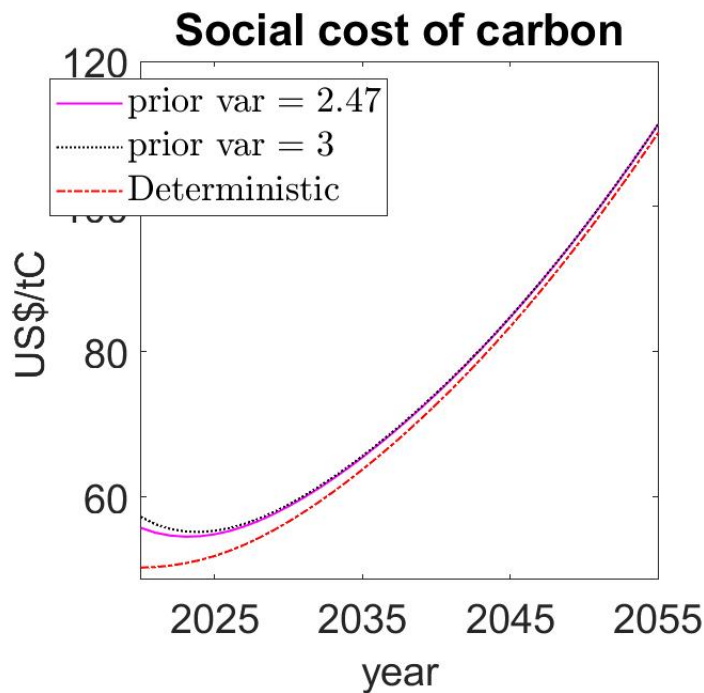


Figure 4: Optimal SCC for different prior distributions over the climate sensitivity. Compares prior variances of 2.47 and 3 to the deterministic SCC.

5.2 Learning

The literature finds that learning is a slow process. Kelly and Kolstad (1999a) find a theoretically expected learning time over climate sensitivity between 90 and 160 years. The damage function in IAMs is the subject of high uncertainty, as every new temperature increase represent unknown territory, a slow learning phase is also reasonable for the convexity of damages. However, improvements in for instance computational power that allows for more comprehensive predictive models, may make the case for a shorter learning period. Therefore an interesting analysis is to vary the speed of learning to see how that affects policy.

In this implementation of a learning model, we've introduced an exogenous falling variance. This does not only have the benefit of reducing the state space, it also allows us to easily set the rate of the declining variance, and hence the speed of learning. This way we can easily compare how fast and slow learning affects optimal policy in this model. The way we do that is by varying the variance over the noise in the model, σ_{ε}^2 . In figure 5 the difference in the SCC from a σ_{ε}^2 of 2 versus 20 is shown. Only the effect for the years up to 2055 is shown as this is the period where the effect of uncertainty is the most apparent.

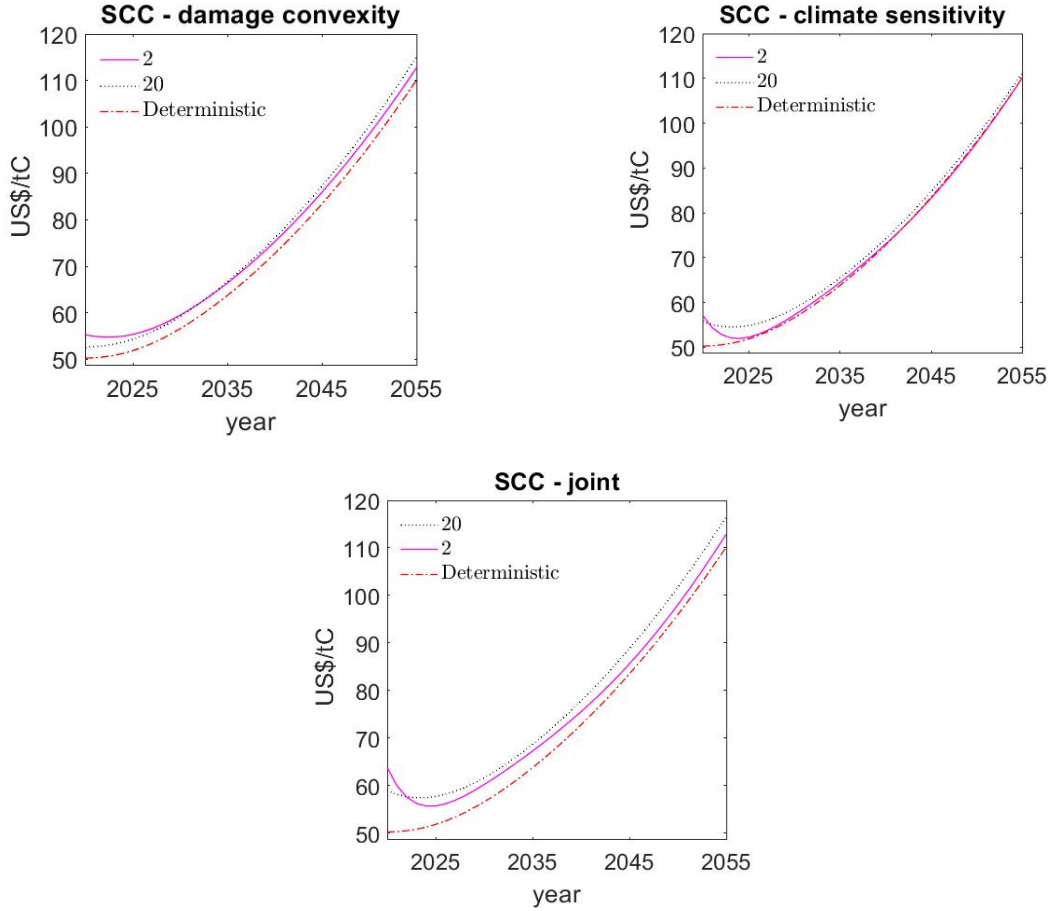


Figure 5: Comparison plots over the SCC. Compares model runs for two different speed of learning. More specifically the plots depicts the difference σ_{ξ}^2 of 2 versus 20 yields compared to the deterministic case.

Year	Deterministic	Joint - 20	Joint - 2	CS - 20	CS - 2	D - 20	D - 2
2020	50.25	58.95	63.89	55.73	57.02	52.62	55.28
2025	51.28	57.49	55.73	54.57	51.97	53.69	54.98
2030	55.41	60.57	58.99	57.63	55.91	57.99	58.4
2070	159.03	167.48	163.67	160.2	159.97	166.09	162.09
2120	492.71	514.68	497.5	494.21	492.26	512.92	498.25
2220	2188.7	2246.1	2197.96	2191.99	2188.19	2243.19	2199.39

Table 4: SCC for plots in figure 5, and extended for years not shown in the graph. Only uncertainty over the climate sensitivity = CS. Only uncertainty over the convexity of damages = D.

Figure 5 displays the resulting SCC for different speed of learning for joint uncertainty

and single uncertainty. We notice that fast learning means a higher initial SCC compared to slow learning in all cases, while over time this dynamic changes, and so the slow learning makes the case for the highest SCC.

To explain the underlying dynamics behind these results we're going to connect these results to how the variance over each of the state variables vary over time. In figure 6 the variance over the state variables governing the mean of climate sensitivity and the mean of damage convexity in each period t is illustrated for the two different speed of learning. Note that it is the conditional variances over the parameters that are displayed, the variance that the decision maker faces at every given point in time.

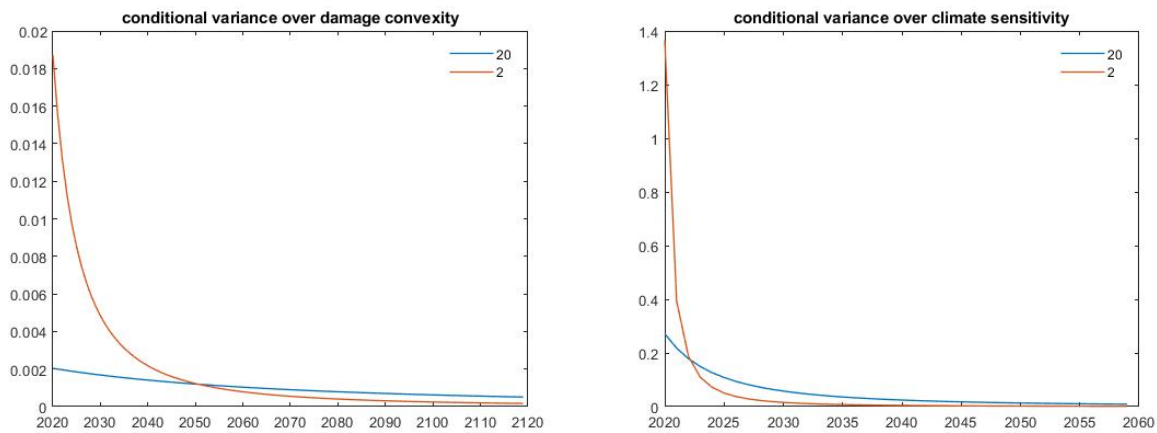


Figure 6: Conditional variance over climate sensitivity and damage convexity. For σ_{ξ}^2 equal 2 and 20.

In both graphs in figure 6, a low σ_{ξ}^2 results in a high initial variance, but which decreases quickly, while a high value of σ_{ξ}^2 results in the opposite. This can explain why we see a higher initial SCC for a low σ_{ξ}^2 in figure 2. Fast learning implies that the decision maker updates his belief over the long-run climate during a short period of time. This results in a substantial updating of the variables during the first few years, something which results in a high initial variance. While for the case of slow learning, the decision maker updates his belief over a longer period of time, which results in a lower initial variance that falls slowly.

From figure 5 we can also see that the fast learning is slower for the convexity of damages than for the climate sensitivity. For climate sensitivity slow learning makes the case for a higher SCC after only 2 years, while for damage convexity uncertainty that is the case after 14 years. This is also what we see in figure 6, there's more updating over a longer time period in the fast learning case compared to the slow, for

the state variable governing damage convexity mean. While for the climate sensitivity slow learning quickly catch up with fast learning.

In the previous section we discussed a hypothesis over why uncertainty only over b_2 has little initial effect on the SCC, while over time it contributes more relative to uncertainty over S_t . This hypothesis was based on the fact that the parameter b_2 has a small effect on damages for low temperatures. Another hypothesis is that we use the same absolute magnitude for σ_ξ^2 that slow learning for both variables. The absolute magnitude of the prior uncertainty and the mean are much lower for damage convexity than for climate sensitivity. Figure 6 shows us that for the same σ_ξ^2 , learning is indeed much slower for damage convexity. This may also explain why uncertainty over climate sensitivity contributes more in the short term, and uncertainty over damage convexity more in the long term.

5.3 Control Rules

To get a more comprehensive understanding of the model we also depict the control rules in figure 7. The control rules let us investigate optimal policy over different realizations of the state variables than along their expected path. Here the control rules for year 2040 and 2240 for the SCC over the climate sensitivity and the convexity of damages are depicted. The other state variables are held fixed at their expected year 2040 and 2240 level.

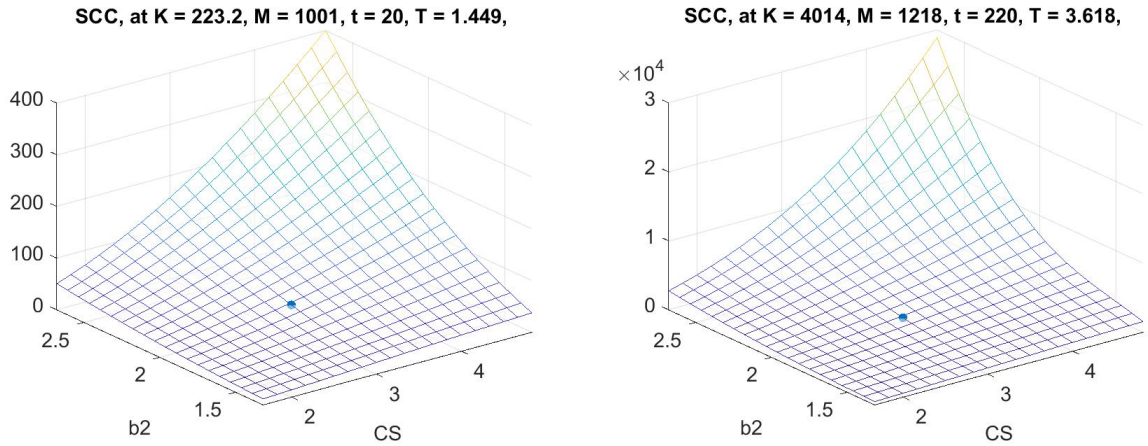


Figure 7: Control rules for the SCC over climate sensitivity (CS) and the damage convexity (b2). The blue dots marks the expected path value for the year 2040 and 2240.

Looking at figure 7 we notice that the curvature of the two years are more or less the

same. For low values of S_t the impact on the SCC from $b_{2,t}$ and vice versa is minimal. That is because for low temperature levels the convexity of the damage function does not contribute much to reducing economic output. For high values of both variables the SCC is much higher. The overall SCC is higher in year 2240, which is due to higher damage realizations from higher temperature levels, and the fact that we discount the future.

In figure 8 the control rules for the joint and single uncertainty runs are compared. We take the control rules for the joint uncertainty run and subtract the control rules for the each of the single uncertainty runs. Hence the graphs show the difference in the SCC between the joint and the single uncertainty runs. When the graphs show a positive value, the SCC for the joint uncertainty run is higher compared to single uncertainty SCC. Here the SCC for different values of the climate sensitivity and damage convexity is shown for years 2040 and 2240. The other state variables are held fixed at their year 2040 and 2240 expected value.

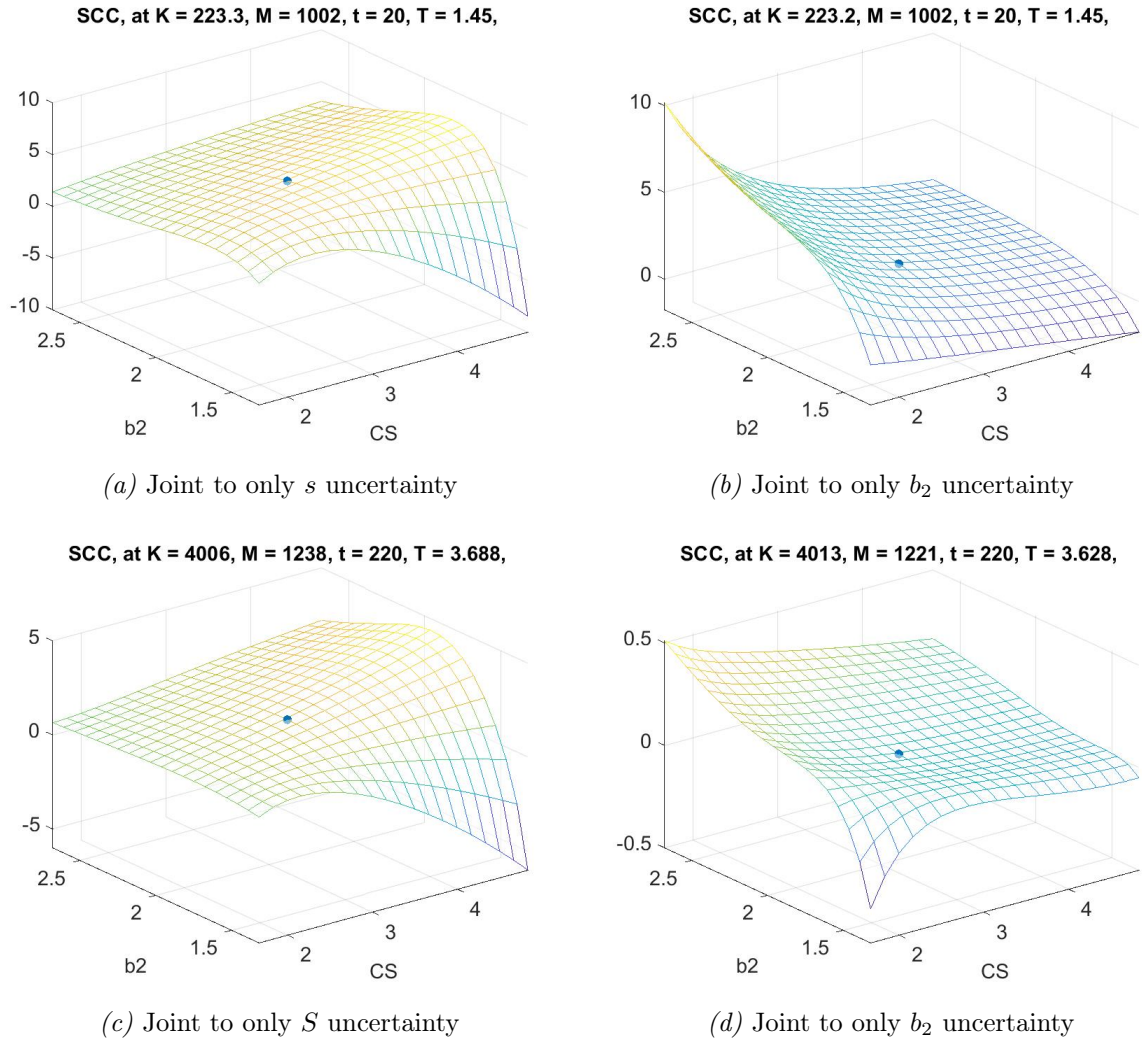


Figure 8: Difference in control rules for joint and single uncertainty runs at year 2040 (figure a and b) and 2240 (figure c and d). Depicts optimal SCC for different values of S and b_2 . The other state variables are held fixed at their 2040 expected value for figure a and b, and at their year 2240 expected value for figure c and d.

Figure 8a and 8c show that only climate sensitivity uncertainty is more dominant for lower values of b_2 . From figure 8b and d we see that only uncertainty over b_2 is more dominant for high values of climate sensitivity. Which is expected as damages are a convex function of temperature.

In figure 8c and 8d we show the same just for year 2240. Figure 8c has the same curvature as 8a. Which indicates that uncertainty only over S_t have the same type of effect in both time periods. The difference is smaller, but that is due to the diminishing effect of uncertainty over time due to learning. Figure 8d differ from figure 8b in curvature. For year 2240 the difference between damage convexity uncertainty

and joint uncertainty is small for all values of S_t . Which is in line with what discussed above, that uncertainty over b_2 becomes more important over time as we experience higher temperatures.

6 Conclusion

In the present study we have implemented uncertainty over two important aspects of climate change; the climate's temperature response to carbon emissions and economic damages from global temperature increase. This is done by implementing an exogenous learning process in a dynamic programming version the widespread IAM, the DICE model.

Our main result shows that the decision maker's response to joint uncertainty is a SCC of 59\$ today, which is 8.5\$ (17.3%) higher than the deterministic SCC. However, this difference decreases quickly over time as uncertainty is resolved.

When considering the two uncertainties separately we find that uncertainty over the climate sensitivity is more dominant in the beginning of the period. This effect is due to the low observed temperature, and a higher prior variance over the climate sensitivity. Over time the uncertainty over the convexity of damages becomes more influential as temperatures increase. The results show that both uncertainties contribute to a higher SCC compared to the deterministic SCC over the simulated period.

Depicting the control rules gave us a more comprehensive understanding over how the interaction between climate sensitivity and damage convexity contributes to the SCC. Here we saw that for low values of both or just one of our two uncertain variables, neither made the case for a high SCC. While for high values of both variables the corresponding SCC increased relatively much.

Inspecting the two different speeds of learning taught us that fast learning implies a higher initial SCC than slow learning. This is due to substantial updating in the beginning of the period, but the effect wears off as the decision maker learns fast. Slow learning made the opposite case.

A limitation in the current study is the intervals employed over the mean of damage convexity, and especially over the climate sensitivity mean. Wider intervals would be beneficial. As there is a trade-off between higher order approximation of the value function and stability, increasing the intervals made us encounter convergence issues, and the present intervals are the highest the model manages to reach at this point. An expansion of the model would be to increase these intervals.

A further expansion of the model would be to employ a right-skewed distribution over the mean of climate sensitivity and damage convexity. There is evidence in the literature that both variables are governed by such distributions (Wagner & Weitzman,

2015).

This study employs expected utility. Within this framework we cannot separate between risk aversion and the propensity to smooth consumption over time. An interesting extension of the model would be to implement Epstein-Zin preferences. Crost and Traeger (2014) show that with uncertainty over the damage function, the policy implications from uncertainty is different under Epstein-Zin preferences as compared to expected utility.

References

- Bellman, R. E. (1957). *Dynamic programming*. Princeton University Press.
- IPCC. (n.d.). *Definition of Terms Used Within the DDC Pages*. Retrieved 04.05.2020, from https://www.ipcc-data.org/guidelines/pages/glossary/glossary_hi.html
- IPCC. (1990). *Climate Change, The IPCC Scientific Assessment*. Cambridge University Press.
- IPCC. (2013). Summary for Policymakers. In *Climate change 2013: The physical science basis. contribution of working group 1 to the fifth assessment report of the intergovernmental panel on climate change*. Intergovernmental Panel on Climate Change.
- JGCRI. (n.d.). *GCAM Model Overview*. Retrieved 31.05.2020, from <http://jgcri.github.io/gcam-doc/overview.html>
- Nobel Media AB 2020. (n.d.). *William D. Nordhaus - Facts - 2018*. Retrieved 31.05.2020, from <https://www.nobelprize.org/prizes/economic-sciences/2018/nordhaus/facts/>
- Charney, J. (1979). *Carbon Dioxide and Climate: A Scientific Assessment*. National Academy of Sciences Press.
- Crost, B., & Traeger, C. P. (2013, May 29). Optimal Climate Policy: Uncertainty versus Monte Carlo. *Economics Letters*, 552–558.
- Crost, B., & Traeger, C. P. (2014, June). Optimal CO2 Mitigation under Damage Risk Valuation. *Nature Climate Change*, 4, 631–636.
- Fitzpatrick, L. G., & Kelly, D. L. (2017, June 1). Probabilistic Stabilization Targets. *Journal of the Association of Environmental and Resource Economists*, 4, 611–657.
- Greenstone, M., Kopits, E., & Wolverton, A. (2013, January). Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation. *Review of Environmental Economics and Policy*, 7, 23–46.
- Howard, P. H., & Sterner, T. (2017, June). Few and Not So Far Between: A Meta-analysis of Climate Damage Estimates. *Environmental and Resource Economics*.
- Hwang, I. C., Reynès, F., & Tol, R. S. (2017, May). The effect of learning on climate policy under fat tailed risk. *Resource and Energy Economics*, 48, 1–18.
- Jensen, S., & Traeger, C. (2016). Pricing Climate Risk. Retrieved 21.05.2020, from https://are.berkeley.edu/~traeger/pdf/Jensen%20Traeger_Pricing%20Climate%20Risk.pdf
- Jensen, S., & Traeger, C. P. (2014, July). Optimal climate change mitigation under long-term growth uncertainty: stochastic integrated assessment and analytic

- findings. *European Economic Review*, 69, 104–125.
- Judd, K. (1998). *Numerical Methods in Economics*. MIT Press.
- Kelly, D. L., & Kolstad, C. D. (1999a). Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control*, 23, 491–518.
- Kelly, D. L., & Kolstad, C. D. (1999b). Integrated Assessment Models for Climate Change Control. *International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues*.
- Kelly, D. L., & Zhuo, T. (2015, July). Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management*, 72, 98–122.
- Knutti, R., Rugenstein, M. A. A., & Hegerl, G. (2017, September 4). Beyond equilibrium climate sensitivity. *Nature Geoscience*, 10, 727–736.
- Lemoine, D., & Rudik, I. (2017, October). Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point. *The Annual Review of Resource Economics*, 9(1), 117–142.
- Miranda, M. J., & Fackler, P. L. (2002). *Applied computational economics and finance*. MIT Press.
- National Oceanic and Atmospheric Administration. (2020, March 5). *Monthly Average Mauna Loa CO₂*. Retrieved 05.04.2020, from <https://www.esrl.noaa.gov/gmd/ccgg/trends/>
- Nordhaus, W. (1992, February). The "DICE" Model: Background and Structure Of a Dynamic Integrated Climate-Economy Model of Global Warming. *Cowles Foundation Discussion Papers 1009*.
- Nordhaus, W. (2007, July 24). *The Challenge of Global Warming: Economic Models and Environmental Policy*. Retrieved 15.04.2020, from http://www.econ.yale.edu/~nordhaus/homepage/OldWebFiles/DICEGAMS/dice_mss_072407_all.pdf
- Nordhaus, W. (2013, October). *DICE 2013R: Introduction and User's Manual* (2nd ed.). Retrieved 15.04.2020, from <https://sites.google.com/site/williamdnordhaus/dice-rice>
- Nordhaus, W. (2017, January 31). Revisiting the social cost of carbon. *Proceedings of the U.S. National Academy of Sciences*.
- Pitesky, M. E., Stackhouse, K. R., & Mitloehner, F. M. (2009). Chapter 1 - Clearing the Air: Livestock's Contribution to Climate Change. *Advances in Agronomy*, 103, 1-40.
- Rudik, I. (n.d.). Optimal Climate Policy When Damages are Unknown. *American Economic Journal: Economic Policy (Forthcoming)*.
- Traeger, C. P. (2014, August). A 4-stated DICE: Quantitatively addressing uncertainty

- effects in climate change. *Environmental and Resource Economics*, 59(1), 1–37.
- Traeger, C. P., & Lemoine, D. (2014, February). Watch Your Step: Optimal Policy in a Tipping Climate. *American Economic Journal*, 6(1), 137–166.
- Wagner, G., & Weitzman, M. L. (2015). *Climate Shock: The Economic Consequences of a Hotter Planet*. Princeton University Press.

A Parameters

Economic parameters		
η	2	Intertemporal consumption smoothing preference
RRA	2	Coefficient of relative Arrow-Pratt risk aversion
b_1	0.0027	Damage coefficient
δ_u	0.015	Pure rate of time preference per year
L_0	6838	In millions initial population.
L_∞	10500	In millions, asymptotic population
g_L^*	0.13449	Rate of convergence to asymptotic population
δ_K	0.1	Depreciation rate of capital per year
κ	0.3	Capital elasticity in production
A_0	0.0067	Initial labor productivity
$g_{A,0}$	0.0235	Initial growth rate of labor productivity
δ_A	0.006	Rate of decline of productivity growth rate per year
σ_0	0.1334	Initial CO2 emissions per unit of output.
$g_{\sigma,0}$	-0.01	Initial rate of decarbonization per year
a_0	1.2613	Cost of backstop in 2013
a_2	2.8	Cost exponent
g_Ψ^*	0.005	Rate of convergence from initial to final backstop cost
δ_σ	-0.1%	Rate of decline of the rate of decarbonization
Climatic parameters		
M_{pre}	588	In GtC, preindustrial stock of CO2 in the atmosphere
$\delta_{M,0}$	-1%	Initial rate of decarbonization
δ_σ	-0.1%	Rate of decrease in growth rate of decarbonization
B_0	0.4202	In GtC, initial CO2 emissions from LUCF
δ_B	-0.04	Growth rate of CO2 emissions from LUCF per year
η_{forc}	3.8	Forcing of CO2 doubling
λ	0.315	Ratio of forcing to temperature increase under CO2 doubling
EF_0	-0.06	External forcing in year 2000
EF_{100}	0.62	External forcing in year 2100 and beyond
σ_{forc}	2.6%	Warming delay, heat capacity atmosphere, annual
σ_{ocean}	0.18%	Parameter governing oceanic temperature feedback, annual
$\sigma_{0,s}^2$	2.47	Prior variance, climate sensitivity
σ_{0,b_2}^2	0.203	Prior variance, convexity of damages
σ_ε^2	2, 20	Variance over the stochastic shock

Table 5: Model parameters

B Comparison to 4 state model

Here we compare our deterministic 6 state model to the 4 stated DICE model (Traeger, 2014). We want to test if the models coincide. The uncertainty is "switched off" in the 6 state model by optimizing the value function without Gauss-Legendre quadrature as explained in section 4.2.2. Over the first years it almost perfectly overlap, while in the end of the time period shown in the graphs, there's only a minor difference between the two.

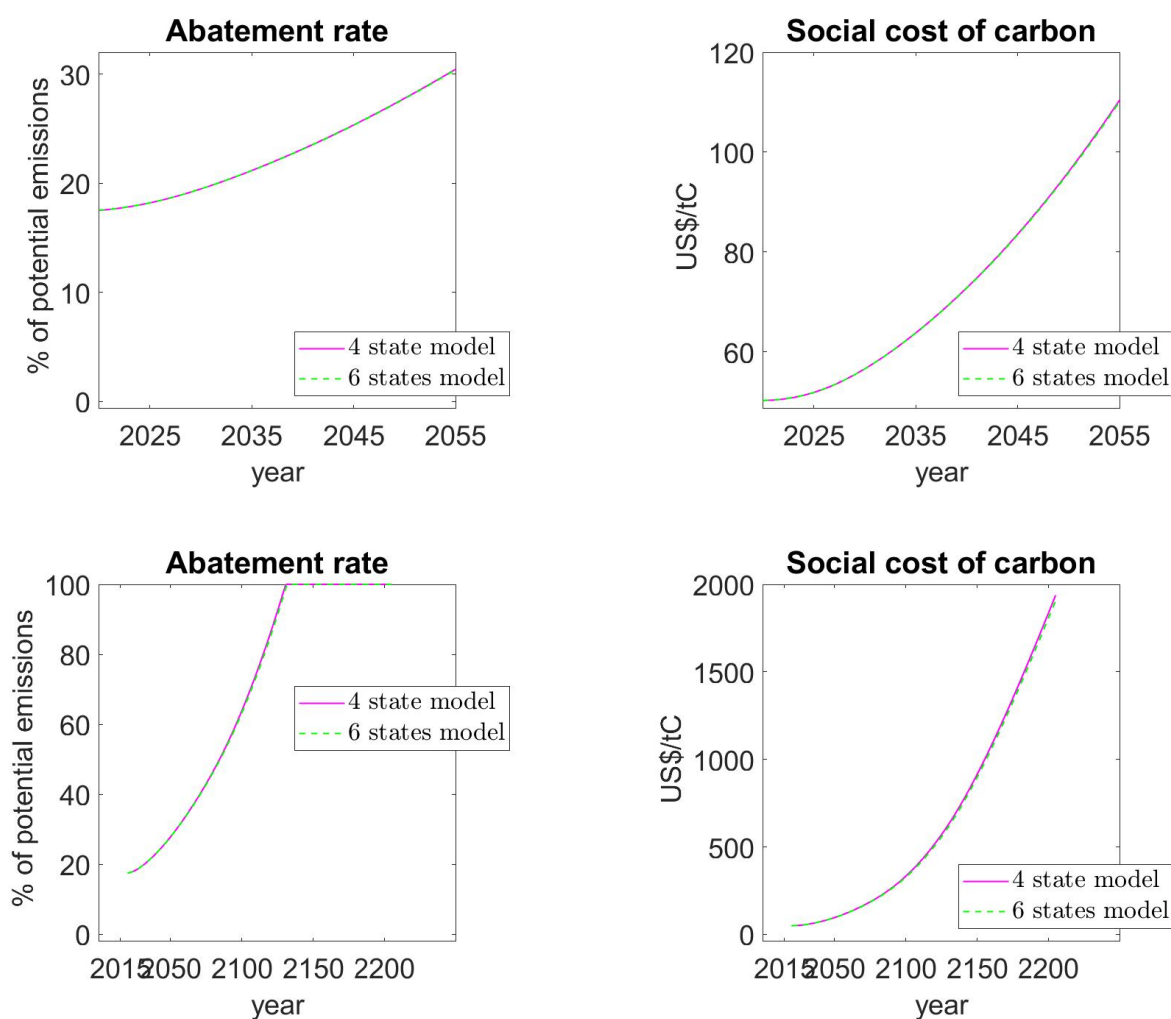


Figure 9: Comparison plots over the SCC and abatement rate. Compares the deterministic 6 state model to the deterministic 4 state model. The top figures shows the time paths simulated for the first 50 years, while the bottom figures shows the same over a 200 years time horizon