Universiteit Utrecht

Faculty of Science

Simulating unknown climate sensitivity in a FAIR-DICE model with SRM

Bachelor Thesis

Fons Jongema

Natuur- en Sterrenkunde

Supervisor:

Dr. Claudia Wieners IMAU

16th of June

Abstract

In this study, we will look at how the uncertainty in climate sensitivity influences policy in a FAIR-DICE climate model with a spread of different values in climate sensitivity. Varying possibilities for the climate sensitivity combined with the possibility for solar radiation management (SRM) as a means of climate change abatement. This way, we hope to gauge the theoretical viability of SRM despite the fact that climate sensitivity is not certain.

Contents

1 Introduction

With the expected rise of global temperature in the coming century, one of the more prevalent questions of climate change is how much anthropogenic greenhouse gases will increase radiative forcing, and thus the surface temperature. One radical way to slow or prevent this is to change the make-up of the atmosphere through geo-engineering, specifically through Sulphur Aerosol Injections, or SIA's for short. Doing this would make the atmosphere more reflective to sunlight, but this technology is still in development, and we do not know all the risks involved. Deploying SIA´s could reduce the damage to the climate, but also exacerbate it. Perhaps opting to reduce our emissions of greenhouse gases will be more beneficial than geo-engineering, or perhaps a mix of the two measures is optimal. To answer this question, we can run an integrated-assessment model that tries to find the optimal course of action. DICE is one of these models: It considers every possible option in a climate model and then proceeds accordingly. But is it possible to assess this without knowing how the climate model will react? We do not know for certain how much the temperature will increase based the policy of DICE. The increase in temperature resulting from a doubling of CO_2 and CO_2 equivalent GHG's is known as the climate sensitivity. It would be extremely convenient to determine this quantity precisely, but this has proven difficult even though our understanding of Earth's climate has increased in the past decades. The current climate sensitivity is estimated to be between 1.5° and 4.5deg per doubling of CO_2. In this study we formulate a way to make a spread of possible values for climate sensitivity while also estimating the likelihood of each of those possibilities. By doing this, we make DICE a model that can estimate the best possible policy with geo-engineering and reduction of emissions, while also accounting for the uncertainty in climate sensitivity.

2 Model

The simulation in matlab is a modified version of the Dynamic Integrated Climate-Economy model, or DICE for short. DICE works through assuming that the world's economy is at the disposal of a single decision-maker, who then tries to find a way to maximize the world's 'welfare' with regards to the future: ignoring climate change will impart damage to the economy, while investing in abatement will cost the economy, but reduce damages. The decision-maker tries to formulate the best policy possible given the constraints given in the simulation.

In DICE, the utility function $U(t)$ (in dollars) represents the consumption of products of the world's population. Access to basic resources for quality of life is what keeps everyone happy, so the integral of utility over time is called the welfare J. In our case, one timestep dt represents 10 years, and with a total of 15 steps, a period of 150 years is simulated. The cumulative welfare over this period will determine the policy. A brief overview of how J is calculated is certainly warranted: In the first timestep, the GDP of the world, Y_0, and the amount of CO_2 in the atmosphere, $M_{at}(0)$ (in Gigatons of CO_2 are initialised. Next, the amount of greenhouse gas emissions and their effect on the economy are calculated.

First, the increase of radiative forcing F due to the increase in CO_2 is considered. This is given by:

$$F_{CO2}(t) = F_{2x} \cdot log\left(\frac{C_{at}(t)}{C_{at}(0)}\right)/log(2) \tag{2.1}$$

Where F_{2x} is the amount of radiative forcing gained from a doubling in CO_2. Of course, $M_{at}(t)$ grows based on the collective emissions $E(t)$ of that timestep. The radiative forcing of greenhouse gases then cause the temperature to increase. Next, this is used to calculate the damage D. Damage is a unitless measure that is later used to reduce the GDP. Many possible 'damage functions' exist, but we use the following one:

$$D = \psi_t \cdot T(t) + \psi_p \cdot P(t)^2 + \psi_c \cdot \left(\frac{M_{at}(t)}{277.88}\right)^2, [3] \tag{2.2}$$

Where $T(t)$ is the temperature increase since the pre-industrial era (P.I.), and $P(t)$ is the increase in precipitation since P.I. The damage scaling constants ψ are:

ψ_t	$1.703 \cdot 10^{-3}$ $°C^{-1}$
ψ_p	0.4 days^{2}mm^{-2}
ψ_c	$3.31 \cdot 10^{-8}$ GtC^{-2}

This amount of damage then reduces the GDP Y:

$$Y(t) = \frac{Y_{gross}(t)}{1 + D(t)} - Costs \tag{2.3}$$

Where Y_{gross} is the size of the economy adjusted by the satisfaction of energy demand, and $Costs$ are the collective investment of the policy maker into abatement and mitigation of climate change.
The final calculation of each timestep is to turn this new Y into the desired Utility function through a rough approximation of the world's economy:

1. $$C(t) = \frac{Y(t)(1 - sr}{L(t)}$$

2. $$W(t) = \int_0^t r(t) \cdot L(t) \cdot \left(\frac{C(t)^{1-\kappa} - 1}{1 - \kappa} - 1 \right) dt$$

Here, C is the consumption of products of each person, sr is the so called 'saving rate': the amount of money the decision-maker pulls from the GDP to mitigate climate change. We keep this rate fixed for all time at 0.25. $L(t)$ is the size of the world's population in millions.
$W(t)$ then, is the 'Welfare': An expression for everyone's quality of life. Here, $r(t)$ is the Impatience of the decision maker: the longer the simulation runs, the lower this prefactor gets. The decision maker generally prioritizes short-term gain more than long term gain. Secondly, η is the 'inequality aversion parameter': This is meant to adjust the Welfare for the inequality in utility: a reliable, constantly growing amount welfare is more valuable than a fluctuating one. With this model for Welfare, DICE then defines the Utility J as the negative of the increase in Welfare.

So, with these conditions, the decision-maker sets out to maximize Welfare and Utility as follows: Two arguments are used to call a function that calculates the utility given the parameters of the climate model: The aforementioned 'saving rate' sr and the so-called 'abatement fraction' μ. This number represents the fraction of CO_2 emissions that the policy maker will avoid for that timestep. Reducing emissions in this way also costs money that is substracted from the GDP. So, running 'evaluate' once corresponds to calculating the utility for one timestep with specific values for sr and μ.
These two inputs and the function for utility are then passed to a Matlab algorithm called 'fmincon()'. This algorithm seeks to minimize the output the 'evaluate' function by varying the saving rate and abatement fraction. Because the utility J is defined as a negative number, this works until the optimal sr and μ are found.
These final values are then used to calculate their effect on the model for the next timestep, and then this process is looped for the amount of timesteps given. It should also be mentioned that for each simulation, the climate model is also run with saving rates and abatement fractions of zero: The policy maker does nothing. This is called the baseline simulation, and it serves to illustrate the consequences of inaction for many variables.

3 Modifications to DICE

For this study, four adjustments to the classic DICE model have been made. I will briefly explain the first three, as they generally improve the improve the accuracy and variety of options in the model, but are not the main focus of this research. The fourth addition introduces a way to create multiple independent values for the climate sensitivy of DICE, and it is this new feature that I used to get my results.

3.1 FAIR-DICE

DICE assumes that the climate has several carbon sinks, mainly the ocean, that can absorb CO_2 from the atmosphere and thus mitigate the increase of temperature through radiative forcing. However, DICE does not focus particularly on the specifics of this absorption process and just assumes that the ocean takes in a certain fraction of the atmospheric CO_2 in each timestep and also gives back a part of its carbon contents to the atmosphere. This approximation done away with in FAIR: the Finite-Amplitude Impulse-Response model for calculating the coupling between carbon sinks and the atmosphere.
This is done by integrating the following differential equation over each timestep:

$$\frac{dR_i}{dt} = a_i \cdot E - \frac{R_i}{\alpha \cdot \tau_i}, \text{ where } i \in \{1, 2, 3, 4\}, [2] \tag{3.1}$$

Here, R_i are the carbon contents of four abstract reservoirs: The findings of FAIR are based on fitting these equations to other climate models. Negative fluctuations in R represent uptake into the atmosphere, positive fluctuations are net depositions from the atmosphere. a_i are the scaling constants that determine how much of the emissions E are absorbed into each reservoir, and the $\alpha \cdot \tau_i$ terms determine how much of the carbon content in the reservoir diffuses back into the atmosphere.
The increase in temperature in FAIR is then calculated as follows:

$$\Delta T_i(t) = q_i \cdot F_{CO2}(t-1) - q_i \cdot F_{CO2}(t-1) \cdot exp\big(-\frac{dt}{d_i}\big), \text{ where i} \in \{1, 2\}. \tag{3.2}$$

Here we take the sum of T_1 and T_2 to be effective increase in surface temperature. The scaling coefficients are $q_i = (0.0659, 0.6684)\,^{\circ}Cm^2W^{-1}$ and $d_i = (239, 4.1)y$
The results of integrating FAIR with DICE are already known to give a more precise representation of the behaviour of these carbon sinks, so I won't go into further detail here.

3.2 Solar Radiation Management

Recently, the possibility of geo-engineering in order to mitigate climate change. One of these possibilities is with so-called SAI's: Sulphur Aerosol Injections into the atmosphere can increase the albedo of the atmosphere, and so reduce radiative forcing and slow the increase in temperature.
Considering this possibility in the DICE-model means that some of the resources given by the saving rate can be spent to start with this Solar Radiation Management, in addition to using it to perform abatement of emissions. The cost of sulfur injections are included in the *Costs* substracted from the GDP. SAI's will persist for only a year in the atmosphere. Thus, great amounts of SAI's must be injected yearly to get a notable effect.
The negative amount of forcing done for a given injection rate of $S(t)$, in megatons of sulfur per year is:

$$F_{srm}(t) = -\eta_s \cdot \alpha_s \cdot exp\bigg(-\frac{\beta_s}{S(t)}^{\gamma_s}\bigg) \tag{3.3}$$

Here, η_s is the estimated effectiveness of the SAI at 74%, α_s is a scaling constant that determines how much forcing is avoided for a given injection rate of sulfur, set to 65 $W \cdot m^{-2}$. β_s is the decay rate for sulphur in the atmosphere at 2246 $Mt \cdot y^{-1}$. Lastly, γ_s is a scaling factor of 0.23 [3]. This quantity is then substracted from the total forcing happening in the baseline DICE-model, thus mitigating the increase in temperature.
However, this theoretical reduction in radiative forcing is not without complication: It is likely that SAI's will not only reduce radiative forcing, but also induce damage to the climate in other ways. Sulphur can damage the ozone layer and make acid rain, which is far from optimal. Still, we assume that SAI's reduce the radiative forcing by enough for it to still be worth the trouble. The amount of damage incurred by sulphur is added to the already existing damage:

$$D_s = \psi_s \cdot S(t)^2 \,\text{with}\, psi_s = 9.25 \cdot 10^{-5} y^2 MtS^{-2} \tag{3.4}$$

The injection rate S is used as the second decision-making variable in DICE: The saving rate determines what fraction of the GDP is used to counteract climate change, but then we must choose whether to spend this money on normal abatement or the possibly risky SAI's. The injection rate S and abatement fraction μ become the two variables that the policy maker can choose.

3.3 Energy model

Next, we added an original system to simulate the means of energy production. In the old DICE, money of the GDP could be spent to achieve a certain abatement μ, costs scaling with $\mu^{2.6}$. This rate would shrink over time, under the assumption that abatement becomes easier as technology advances. So, the decision maker would often wait until later steps before investing in abatement, profiting from the free discount.
This was of course unrealistic, and so we have implemented a new way to do abatement. In the new energy model, a more involved way of achieving abatement is installed. In each timestep, the amount of emissions E is generated by power plants. The initial amount of emissions of the old DICE is attributed to the polluting power plants in the world. So these 'brown' plants initially contribute 85% of the total power production, the remainder is done by green plants that give zero emissions. Then, for each time step, the energy demand must be met or the GDP will be reduced by the power deficit, lowering the utility. The decision maker almost always avoids this by investing enough resources for the energy demand: the fossil fuels for the brown plants, the building of new green and brown plants, and the replacing of old plants is all accounted for.
Abatement is still by definition the fraction of CO_2 avoided, so consequently, this becomes the fraction of energy generated by green plants over the total amount of energy production. The abatement fraction μ is replaced by the amount of newly built green plants as one of the variables for the policy-maker. In addition to this, the price of building green plants becomes lower by a certain 'learning factor' for each green plant built, while the price of running brown plants increases as fossil fuels become more scarce. In this way, a more consistent way and realistic way of simulating the costs of abatement is achieved.

3.4 Variability in Climate Sensitivity

Last but not least, we introduced a statistical way to give a spread of values for climate sensitivity. To start, we define the real climate sensitivity as the so-called Equilibrium Climate Sensitivity or ECS: This represents the temperature increase per doubling of CO_2 after waiting infinitely long. The climate is of course never in equilibrium, but in theory a certain level of warming will be reached after a very long time. To calculate ECS, we use two other variables: The Realised Warming Fraction, RWF, and the Transient Climate Response, TCR.
The transient climate response is defined as the increase in temperature per doubling of atmospheric CO_2 after 70 years, while the climate is not yet in equilibrium. That is, the change in temperature is not assumed to be final. RWF is defined as the amount of temperature increase that is reached after a certain amount of time divided by the maximum expected increase. Thus, it always lies between 0 and 1. For the purposes of our study, we can assume that RWF and TCR are statistically independent: A change in the one does not influence the other, and vice versa [1]. With this in tow, we now use a combination of these two quantities to find the ECS:

$$ECS = \frac{TCR}{RWF} = F_{2x}(q_1 + q_2) \tag{3.5}$$

Here, q_1 and q_2 are the variables that FAIR then uses in its climate model. q_1 roughly equates to the amount of time needed for the deep ocean to reach thermal equilibrium, and q_2 is the same but for the upper ocean [1]. These two values can be solved for in Matlab and then integrated into the FAIR-model to accurately simulate the effect of a certain value for ECS on the climate.
With this approach to calculating climate sensitivity, we can confront the decision-maker in ordinary DICE with a new problem: What to do if climate sensitivity is uncertain?
First, we define RWF as a normal probability density function in matlab, with a mean $\mu_{RWF} = 0.6$ and a spread of $\sigma_{RWF} = 0.09$. This way, its pdf is normalized and ranges from about 0.3 to 0.9. Second, we define TCR as a log-normal distribution with $\mu_{TCR} = 0.458$ and $\sigma_{TCR} = 0.278$. This is also a normalised pdf, albeit more slanted to the left.
Next, we calculate the cumulative density functions of both of these, and then sample from these CDF's over intervals of equal probability and take the mean value of RWF or TCR from that interval. Two arrays, each with the length of the sample size N, are obtained. The values range from low to high, but all have the probability of $1/N$. In the figures below, a sample of the RWF and TCR is made for every Nth part of the surface under the slope.

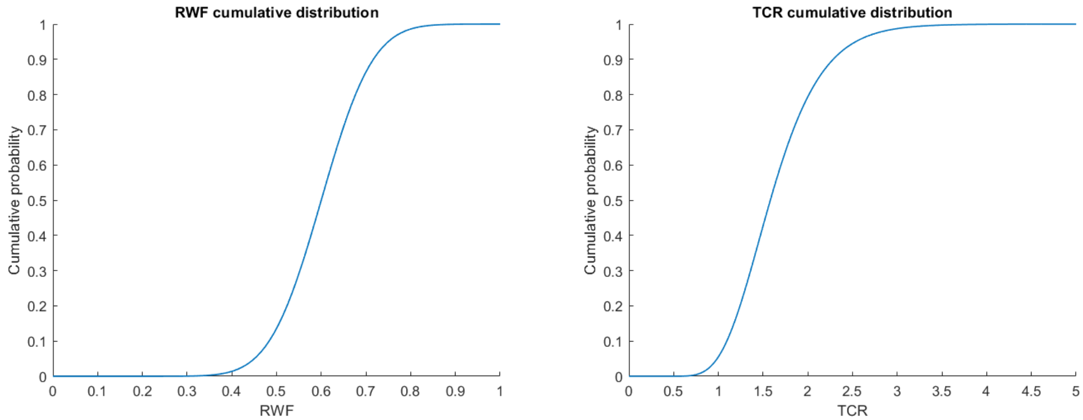

(a) RWF CDF

(b) TCR CDF

Figure 1: The two cumulative distribution functions

From here on, ECS, q_1 and q_2 are calculated with these arrays. This convolutes the two arrays, and in total N^2 values for ECS are obtained. The model then runs with each of these values approximate a policy for an uncertain climate sensitivity. This policy weighs the welfare of each possible outcome equally, as they are all as likely as N^-2. This expected welfare is then maximized for each timestep. So, the policy-maker never finds out what the ECS truly is, but can still plan ahead despite the uncertainty.

4 Results

4.1 Runs

A total of four simulations were done, each with varying values for the RWF and TCR pdf's. The first run is done with a sample size of $N = 5$ to utilize the probabilistic nature of ECS. The next three runs were made with $N = 1$, to compare the differences between uncertainty and a single value for ECS. In these cases, only one value for TCR and RWF is taken into consideration. The spread doesn't matter in that case. The second run uses the reference values for TCR and RWF, the means of both the PDF's. The third run uses the lowest possible ECS possible in the $N = 5$ scenario, which is the most optimistic case. In the final scenario, the highest possible ECS of the $N = 5$ is used to simulate the worst-case scenario. The values for each run are the following:

	μ_{RWF}	σ_{RWF}	μ_{TCR}	σ_{TCR}
1: N = 5	0.60	0.909	0.458	0.287

	TCR	RWF	ECS
2: reference	1.58	0.60	2.63
3: optimal	1.06	0.727	1.46
4: worst case	2.39	0.47	5.09

4.2 Results

For the results, I will show and compare the variables most relevant to our questions: : these are temperature, damage, CO_2 in the atmosphere, and abatement, and the amount of SAI's.
For temperature and damage, a total of N^2 curves are calculated each run. As each sample for ECS is equally probable, I took the mean of all these cases to be the most accurate, and also plotted the lowest and highest values for clarity.

First of all, the temperature. It can be seen that RWF and TCR have quite a large effect: a difference from 3 to 8 degrees increase in temperature is possible in the baseline scenario of $N = 5$. The mean appears to be 4.5, tough. As expected, temperature rises much more strongly in any baseline scenario, but gets the highest with high ECS. Note that in the $N = 5$ uncertainty case, the three baseline curves each roughly correspond with the optimal, reference, and worst-case baselines of $N = 1$, as was the intention with the sampling of TCR and RWF.

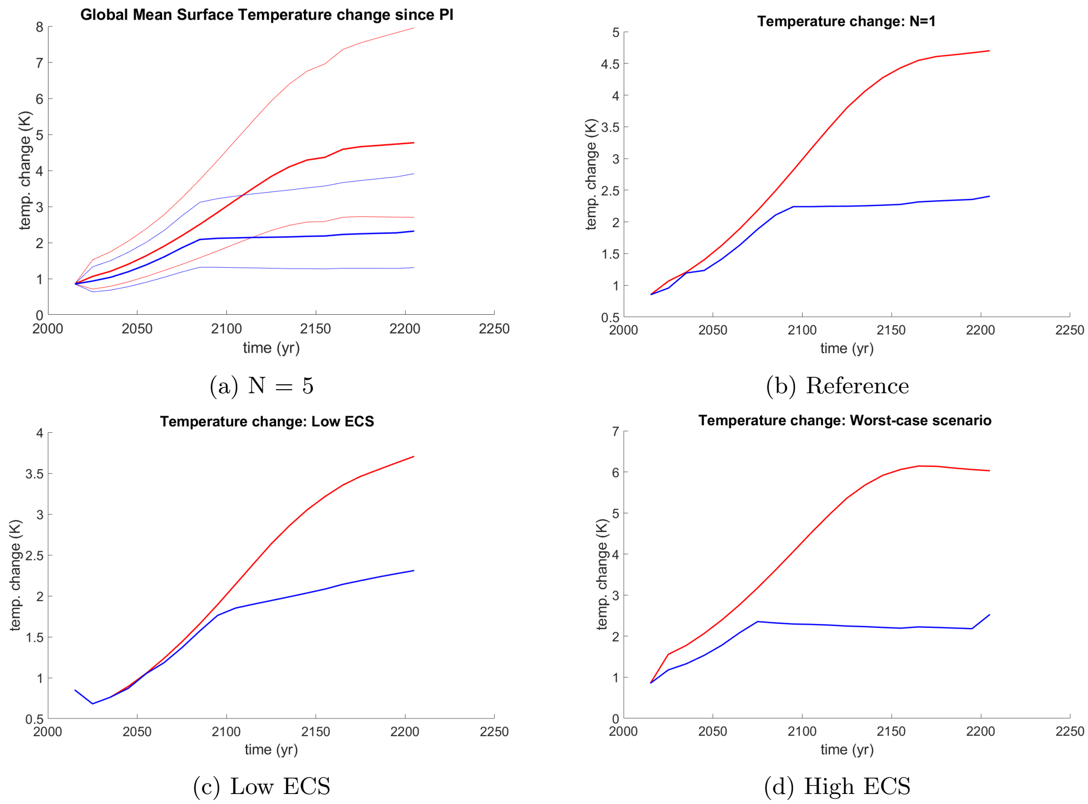

(a) N = 5

(b) Reference

(c) Low ECS

(d) High ECS

Figure 2: Temperature change over multiple runs

Next the damages associated with this rise in temperature in each possible scenario. Damage is roughly equal in the reference and optimal scenarios, while in the worst-case scenario, damage first rises only for it to be reduced later. In the scenario with uncertainty, damage rises at first, but plateaus after the full transition to green energy.

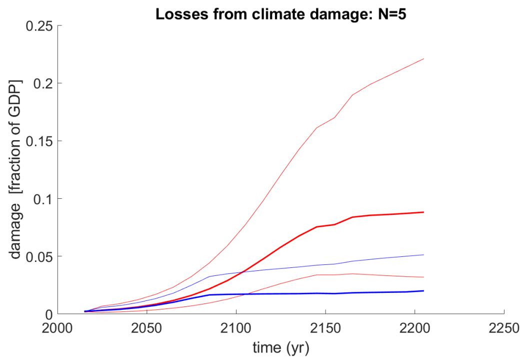

(a) N = 5

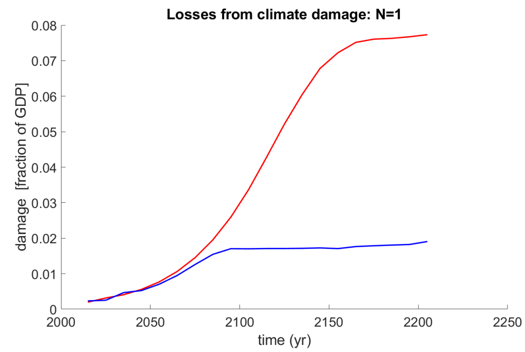

(b) Reference

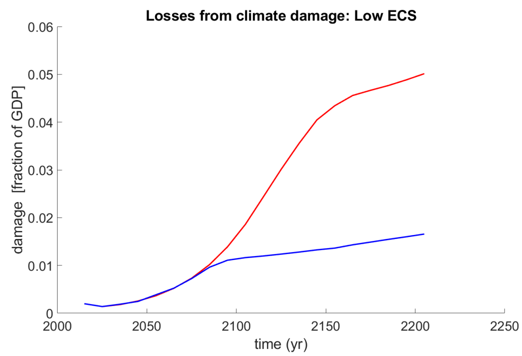

(c) Low ECS

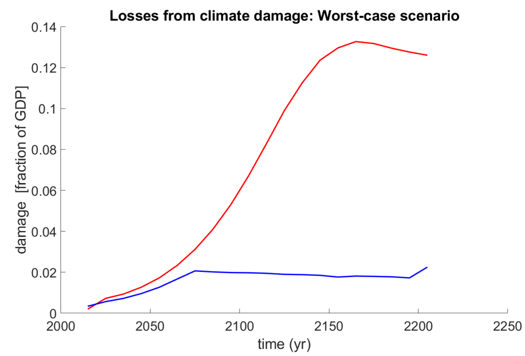

(d) High ECS

Figure 3: Damages over multiple runs

The CO_2 associated with each run is as follows. Note that in the baseline, the carbon content is mostly the same in each case. With the $N = 5$ scenario, again we see that the policy maker takes a balanced approach. In the most optimistic scenario, carbon is mitigated as late as possible. In the reference case, the peak of carbon content allowed by the policy is earlier. In the worst-case scenario the policymaker correctly recognizes how much climate damage is exacerbated, and adjusts for this by mitigating carbon emissions sooner.

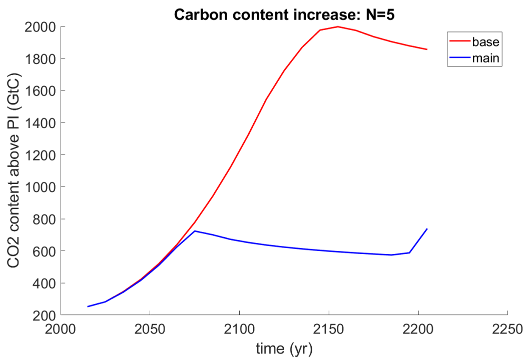

(a) N = 5

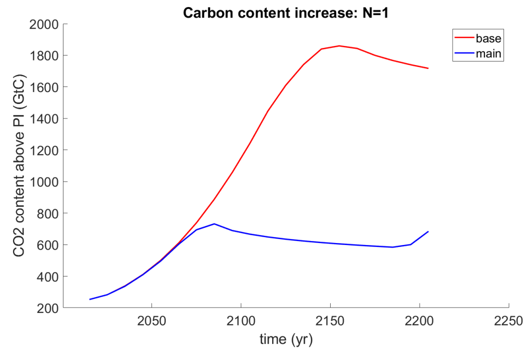

(b) Reference

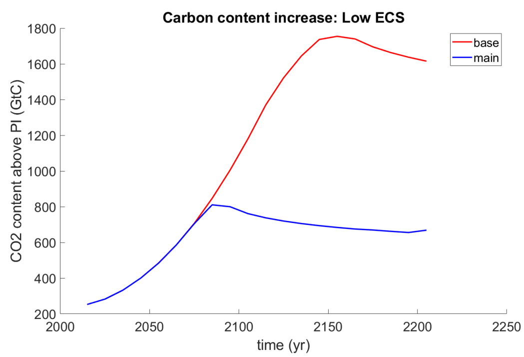

(c) Low ECS

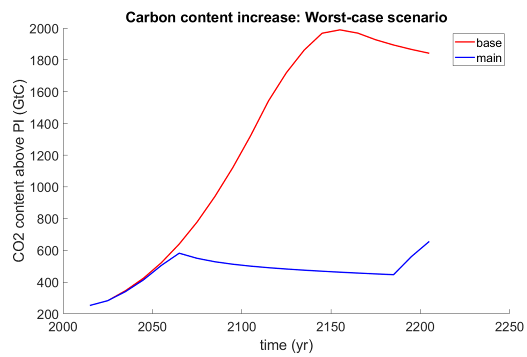

(d) High ECS

Figure 4: CO_2 in the atmosphere

The abatement that the policy maker decides on. Note that in the $N = 5$ run, the policy maker immediately starts enacting an abatement policy around 2050, earlier than in the optimal or reference scenario. This means that the policy-maker acts in a risk-averse way in the case of uncertain climate sensitivity, likely because the high-ECS possibility incurs so much damage. Note that the time it takes to completely transition to green energy is about 20 years in each scenario, the only difference is in when the policy-maker starts.

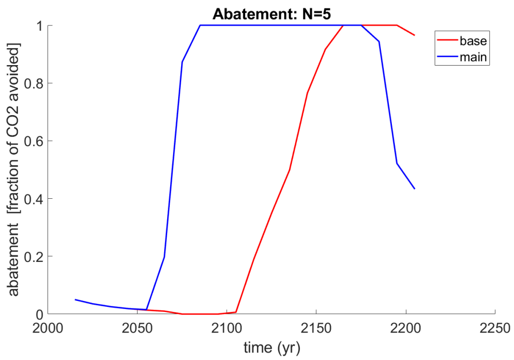

(a) reference values

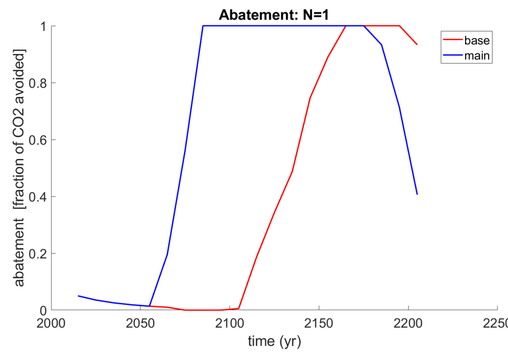

(b) N = 5

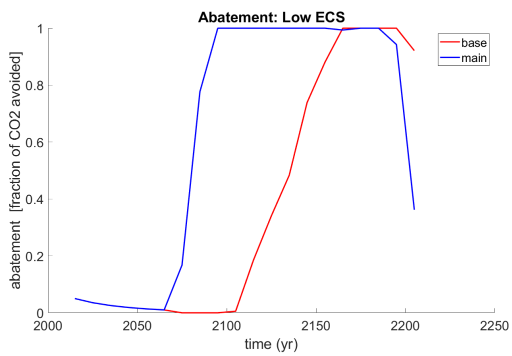

(c) Reference

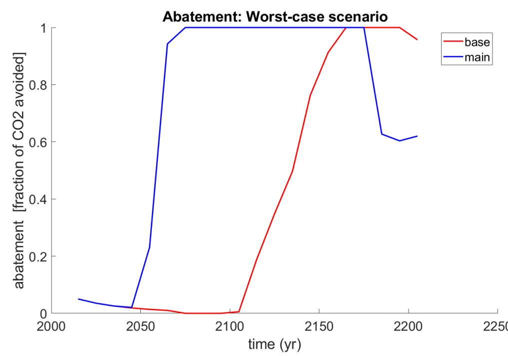

(d) Low ECS

Figure 5: Abatement

And lastly, the amount of SAI's used in each scenario. Again, the same trend is notable: While an increase in the use of SAI's from the optimal to the worst-case scenario, with the reference scenario in between. All the while, The policy for SAI in the $N = 5$ scenario is slightly higher than the $N = 1$ reference scenario, due to the risk-averse quality of this approach.

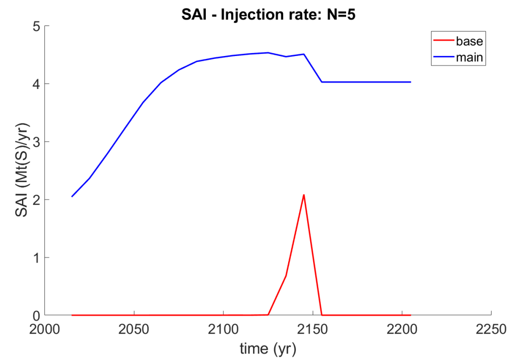

(a) N = 5

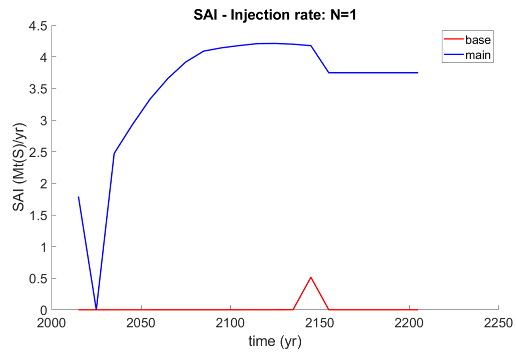

(b) Reference

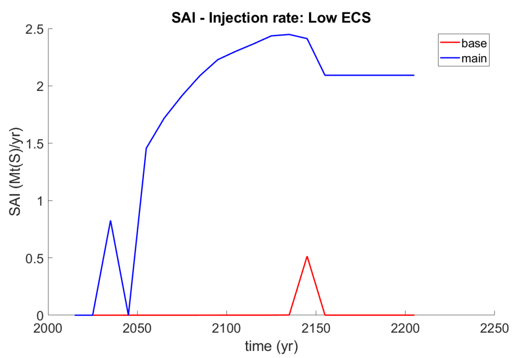

(c) Low ECS

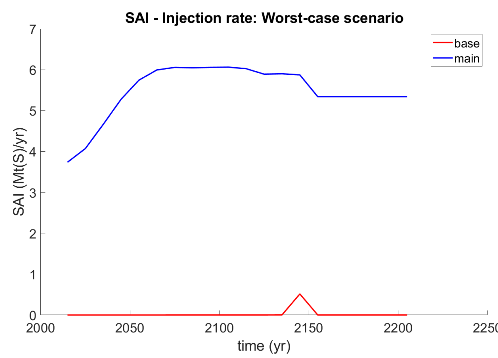

(d) High ECS

Figure 6: SAI

5 Conclusion

In conclusion, we have run a climate model with a spread of values for climate sensitivity, and in this scenario, the optimal policy chosen is an extremely cautious one. A mix of both Solar Radiation Management and normal CO_2 abatement is relied upon for this. In all scenarios, a situation where we transition completely to sustainable energy is eventually reached, the harsher the climate sensitivity, the sooner. Likewise, in all four cases SAI's are also relied upon, but slightly more in the uncertain scenario. The policy-maker seems to mitigate climate change even more strongly if it is given a range of equally likely, but strongly different climate sensitivities to gauge the outcome of its actions. With this method a varying model for climate sensitivity can be used to dynamically inform the policy in one run of an integrated assessment model. As a final remark, it can be gleaned from these results that the mean of all these values for ECS result in a maximum temperature increase of 4.5 degrees, which is very much in line with current predictions. So, it can be assumed that our current understanding of the value of climate sensitivity, while it is impossible to measure directly, is probably somewhat accurate.

6 Discussion

As always, many things about this study can be improved in the future. While the model is stable and rarely crashes because of something unforeseen, it is sometimes slow to run, especially as sampling size N and the amount of timesteps increase. This exacerbates the fact there are many more variables in the climate model that I did not touch on that can be tweaked in order to alter the results. Many of these variables are of course quite abstract, and there is little to no way to verify them that we are simulating something in the future. For example, Solar Radiation Management is currently a concept that is still being developed, so its parameters can only be put into a simulation with some reasonable guesswork.

It is also remarkable that in each scenario, the policy-maker is able to mitigate all of climate change rather quickly and transitions to green energy. The only difference is when the policy-maker starts with abatement. It is strange that this is not done immediately, but the policy-maker instead waits and bides its time. One can wonder why ending climate change as soon as possible does not maximize the welfare in any realistic scenario, and so I find it most likely that this is some unforeseen quirk in the code.

This leads into the next point of unrealistic results. Because DICE is an Integrated Assessment Models wherein all of humanity's resources are at the command of a single policy-maker, some outcomes future policies seem quite impossible. In all scenarios, the policy implemented to drive back temperature change is so effective that temperature increase is stopped before 2050, paired with an abatement that reaches 1 around the same time. This is not necessarily false: we are not in a world were all of our resources can be swiftly expended to this cause, but is doubtful that such a positive outcome -complete reversal of global warming - could be achieved even if that were the case. Many of the results of IAMs should be viewed this in mind.

References

[1] Richard J. Millar, Zebedee R. Nicholls, Pierre Friedlingstein, Myles R. Allen: A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emission, from: Atmosphere Chemistry Physics, 2017, pages 7213-7228.

[2] Timm Faulwasser, Robin Nydestedt, Christopher M. Kellett, Steven R. Weller: Towards a FAIR-DICE IAM: Combining DICE and FAIR-models from: IFAC Papers Online, 2018, pages 126-131.

[3] Koen G. Helwegen, Claudia E. Wieners, Jason E. Frank, Henk A. Dijkstra, Complementing CO2 emission reduction by solar radiation management might strongly enhance future welfare, from: Earth System Dynamics, june 2019, pages 453-472.