Utrecht University

Graduate School of Natural Sciences

Institute for Marine and Atmospheric research Utrecht

Optimal Policy Under Uncertain Climate Sensitivity in an Agent-Based Model

First Supervisor:
Dr. Claudia Wieners

Author:
Ben Romero-Wilcock

Second Supervisor:
Prof. dr. ir. H.A. (Henk) Dijkstra

MSc Climate Physics

December 31st, 2022

Abstract

This report aims to answer two primary research questions: how aggressive policymakers should be in their approach to mitigation while the Equilibrium Climate Sensitivity (ECS) is still largely uncertain, and whether policymakers should adapt their strategy as understanding of the ECS evolves. An agent-based integrated assessment model, the DSK model, is modified to incorporate adaptive policymakers, who learn the climate sensitivity as the global mean surface temperature increases and update their strategy accordingly. The learning process is a combination of Bayesian inference and externally imposed probability distribution functions, which aim to simulate developments in climate science. The outcomes, in global warming and unemployment, seen under adaptive policymakers are compared with the outcomes under non-adaptive policymakers, who maintain the same approach as temperature increases. Risk-neutral policymakers, concerned with the expected ECS are compared with risk-averse policymakers, who are concerned with the 99^{th} percentile value. These four policymakers are compared under two policies: a carbon tax with no accompanying policy, and a carbon tax, 50% of the revenue of which is used to fund the building of renewable energy sources. It is concluded that only risk-averse policymaking is effective, where meeting the Paris Climate Agreement goals are concerned. Among risk-averse policymakers, when the results are aggregated across different ECSs according to their current estimated probability, non-adaptive policymaking achieves the greatest climate mitigation. However it also leads to higher expected unemployment, even under the second policy. While which of these policymaking strategies is preferable may be debatable, it is clear that it is the choice of policy that has the most impact, in terms of both unemployment and climate change.

Acknowledgements

The greatest part of the thanks for this project has to go to my primary supervisor, Dr. Claudia Wieners, without whose guidance and advice I would not have been able to complete half of the analysis in this report. It has been a joy working under her supervision, with our (almost) weekly meetings being both stimulating and a great source of motivation to keep working on the project and thinking about new directions the work could be taken in. I would also like to thank Prof. Dr. Henk Dijkstra for agreeing to be my second supervisor. A great deal of thanks go also to Anke, Andrea and Dr. Anna von der Heydt; our meetings with Claudia discussing our work on the DSK model have been a very enjoyable addition to this thesis project. I would like to thank my flatmates and friends in Utrecht for providing a great environment to work in, not to mention making my time in the Netherlands something I will cherish. I would be remiss at this point not to thank Pierre, particularly for our big walks back in Brussels. Mark and Dan I have to thank for allowing me to pick their brains on this and related topics. Naturally, a big thank you goes to my parents, not least for their proofreading and patience, and lastly but very importantly to Ada, whose support throughout this project was indispensable, not to mention very nice.

List of Commonly Used Acronyms

IPCC	Intergovernmental Panel on Climate Change
IAM	Integrated Assessment Model
DSK	Dystopian Schumpeter Meeting Keynes (Model)
DICE	Dynamic Integrated Climate-Economy (Model)
C-ROADS	Climate Rapid Overview And Decision Support (Model)
FAIR	Finite Amplitude Impulse Response (Model)
ECS	Equilibrium Climate Sensitivity
TCR	Transient Climate Response
PDF	Probability Density Function
GMST	Global Mean Surface Temperature
CESM2	Community Earth System Model Version 2

Contents

Chapter 1

Introduction

As a developing global phenomenon and topic of research, anthropogenic climate change is increasingly important: it has gone from a largely theoretical concern 50 years ago to a key factor in unprecedented floods effecting over 30 million people [1], wildfires north of the Arctic circle [2], and the subject of several years of school strikes [3]. It has been noted that there has recently been an increase in concern about climate change, particularly since the publication of the Intergovenmental Panel on Climate Change (IPCC)'s 'Global Warming of 1.5°C' in October 2018 [4]. This is reflected in the language employed by news media to describe climate change: since 2019, for instance, the UK's Guardian newspaper has been followed by Germany's Der Spiegel, Poland's Gazeta Wyborcza and the Spanish-speaking EFE and Noticias Telemundo, among others, in favouring the terms 'climate emergency' and 'climate crisis' over 'climate change' [5]. This heightened sense of urgency is well-founded: in October 2022, the United Nations Environment Program (UNEP) published their Emissions Gap Report 2022, which noted that 'incremental change' was no longer sufficient [6]. UNEP affirms that 'broad-based economy-wide transformations are required to avoid closing the window of opportunity to limit global warming to [the Paris agreement goals of] well below 2°C, preferably 1.5°C' [6].

Economists working in the fields of ecological and environmental economics have developed differing solutions on just how labour and resources might be mobilised on such an 'economy-wide' scale to effectively mitigate climate change. For some, the climate crisis amounts to a market failure - perhaps 'the greatest example of market failure we have ever seen' - which can be corrected through a combination of carbon pricing, stimuli and regulation [7]. A carbon tax, perhaps the simplest form of carbon pricing, consists of a tax levied on the emittors of carbon-dioxide-equivalent (CO_2e) gas, with a certain value per tonne of CO_2e emitted, with the goal of reducing demand for fossil fuels and encouraging the pursuit of alternatives, thus reshaping the market to reflect the negative externality of global warming [8]. Others believe that more stridently interventionist measures will be required, such as a steadfast commitment by governments to maintain full employment, the mobilisation of a 'carbon army' of workers to build green infrastructure and a broader project of pursuing public ownership of energy infrastructure and financial institutions [9]. Another viewpoint among some environmental and ecological economists is that any effective strategy for climate mitigation would have to be agnostic about - if not completely run counter to - economic growth itself, potentially requiring a conscious effort to free the economy of its perceived growth-dependence [10, 11]. It should be noted that these points of view are not strictly mutually exclusive [12]. To help arrive at - and defend - these conclusions, environmentally-concerned economists employ a range of computational and analytical tools, one of which, of particular interest to this work, is integrated assessment modelling.

1.1 Integrated Assessment Modelling

While the strict definition of what integrated assessment models (IAMs) consist of is not perfectly uniform among environmental economists [13], a broad definition can be developed, based on that in Ackerman et al. (2009) [14]. IAMs are multi-equation computational models combining climate simulations fitted to General Circulation Models (GCMs) with economic models to assess the benefits and costs of different climate policy options. First developed by William Nordhaus through his Dynamic Integrated Climate-Economy (DICE) model in 1992, some IAMs offer the ability to optimise policy according to the quantity of economic welfare. This is achieved through the use of a 'damage function', which determines the amount of damage inflicted by global warming on the modelled economy's productivity, and hence welfare, which is defined so that it is related to GDP through consumption [15].

One relatively recent development in IAMs is the incorporation of the technique of agent-based modelling into their economic component [16]. Agent-Based Models (ABMs) are models which seek to 'replicate the known characteristics and behaviour(s) of real-world system[s]' - and aid in investigating the dynamics of such systems under different conditions - by modelling elements of these systems as autonomous agents [17]. Benefits of agent-based modelling include the fact that the model's agents often map intuitively to components of the system being investigated, potentially making the assumptions upon which the model rests more easily communicable, and the fact that heterogeneous agents can be explicitly modelled, capturing complexities which would otherwise have to be parametrised in some way if a top-down modelling approach had been used [17]. In the context of economic models, non-equilibrium effects, such as business cycles, can be captured without the need for parametrisations, since they may emerge from the interactions of the explicitly modelled individual agent-firms [16]. The Dystopian Schumpeter-Meeting-Keynes model, hereafter abbreviated as the DSK model, is one such agent-based IAM; its structure will be expounded in chapter **??** [16]. The DSK model is not the only agent-based IAM, with at least four other examples existing to the author's knowledge at the time of writing [18].

IAM's have been criticised, with one complaint being the great sensitivity of many IAMs to arguably arbitrarily chosen parameters, such as the rate at which future welfare is 'discounted', relative to the present. The damage function is also a major source of divergence between the results of different studies involving IAMs, and has consequently been criticised along similar lines [19]. It should be noted that the DSK model, in the form used and adapted in this work, makes use of neither a damage function nor discounting, and is thus not affected by these criticisms. Nonetheless, criticisms that could apply to the DSK model include the lack of consideration of 'tail risks' of potentially catastrophically high global warming, and of the large degree of uncertainty which still surrounds the Earth's climate sensitivity [19].

1.2 Climate Sensitivity and IAMs

The equilibrium climate sensitivity (ECS) - the chosen measure of climate sensitivity in this work - is defined as the equilibrium change in global mean surface temperature following a doubling of CO_2 concentrations from pre-industrial conditions [20]. At the time of writing, the IPCC's latest report on the 'physical science basis' of climate change states with high confidence that there is at least a 66% chance that the ECS lies between 2.5 and 4°C/doubling CO_2, placing their best estimate at an ECS of 3°C/doubling CO_2 [21]. Similarly, a 2020 estimation of the probability distribution of ECS by Sherwood et al., used extensively in this work, gives a 66% range of 2.3-4.5°K/doubling CO_2 [22]. Considering that by 2019, atmospheric CO_2 concentrations had already increased by 47% since 1750 [21], and an estimated 20 million people globally would be subject to heat stress exceeding the survivability threshold with just 2.5°C of warming [23], the prospect of an ECS at the upper end of this range - or even beyond it - is major cause for concern, and is deserving of attention by studies that make use of IAMs.

The uncertainty surrounding the ECS formed a large part of economist Martin Weitzman's criticism of the application of cost-benefit analysis to the problem of climate mitigation. Weitzman used the example of a climate sensitivity probability density function (PDF) with a fat tail to argue that such fat-tail probabilities with negative consequences can outweigh the discount factor used, to render cost-benefit analysis practically unusable [24]. While research on the effects of different possible climate sensitivities on IAM outcomes has been conducted since Weitzman's criticism [25], it is a comment by William Nordhaus in his response to Weitzman's original criticism that this work uses as a prompt to pose its primary research questions:

> 'This means that we can learn [the ECS], and then act when we learn, and perhaps even do some geoengineering while we learn some more or get our abatement policies or low-carbon technologies in place.'[26]

This comment, and the assumptions on which it appears to lie, raises several questions. First among these is what the strategy should be when the ECS is still uncertain: when most of the learning is yet to be done. As our conceptual ECS PDF evolves, would it prove wiser to focus only on the expected ECS, or should a conscientious policymaker be more preoccupied with the long tail of the PDF? Second, and perhaps more fundamentally, are there any genuine benefits to changing strategy as the knowledge evolves? As mentioned, there appears to be broad consensus that the window of opportunity for the implementation of policies that could feasibly put global emissions on a pathway to stabilising temperature change below 2 degrees is rapidly fading [6]. The question is thus raised of whether we are already at the stage where the policy that needs to be implemented is the most stringent one that is feasible - in this light, learning that the ECS is higher or lower than previously expected may not necessarily radically alter the policies that might be the most favourable to this end. Finally, there is a broader question of how these policies may be reached without the explicit use of cost-benefit analysis. As noted above, the model used in this work does not make use of a damage function; this, combined with the model's greater complexity and thus computational run-time, renders the optimisation of policy through cost-benefit analysis impractical. While this might complicate the comparison of the outcomes of different policies, it provides certain opportunities for qualitative analysis, as it becomes necessary to analyse in some detail the behaviour of different indicators under different policies.

1.3 Research Questions

In sum, then, the research questions this work will be concerned with can be stated as follows:

1. How aggressive should policymakers be in their approach to mitigation while the Equilibrium Climate Sensitivity is still largely uncertain?
2. Should policymakers adapt their strategy as understanding of the ECS evolves?

It is hoped that the method for policy evaluation that this work follows, combining quantitative prediction and aggregation over different possible ECSs with more qualitative comparison and discussion, will to some extent satisfy the above-mentioned desire to understand how appropriate policies might be selected without relying on explicit cost-benefit analysis.

Chapter 2

The DSK Model

What follows in this chapter is a brief overview of the DSK model and the way its climate module has been modified and calibrated to have the desired initial conditions for this research project.

2.1 Overview

The version of the DSK model used in this work is based on an udpated version of that outlined in Lamperti et al.'s initial presentation of the DSK model [16]. A schematic of the main features of the model is presented in fig. (2.1). The economic component of the model is centred around two sectors, modelled following an agent-based approach in which individual firms are modelled as agents interacting in an imperfect market, which together make up a stylised industrial sector. Sector 1, the capital goods sector, comprises 50 firms which produce the machines which are sold to and used by the 200 firms which make up the consumption goods sector. Consumption goods firms produce a homogeneous consumer good which is sold to households, which are modelled as an aggregated mass of labourers and consumers. The model features a banking sector comprised of private banks which interact with a central bank and the government. Banks play the role of buying government bonds and supplying credit to consumption good firms. Crucial to this work is the energy sector, which consists of a monopoly energy provider which generates electricity from a combination of brown and green plants. Brown plants represent fossil-fuel based energy production, producing CO_2 emissions proportional to the amount of electricity generated, while the green plants are an idealised amalgam of renewable energy sources and produce no emissions. Conceptually, the economic component of the DSK model might be thought of as a model of a small, statistically representative country, rather than an explicit model of the global economy. This distinction is made due to the small number of firms and small working-age population, which is held constant at 250,000 people.

CO_2 can be emitted from two sources in the DSK model. These are the electricity firm, as previously stated, and sector 1 firms, which can source their electricity from the electricity firm or generate it themselves by burning fuel: crucially, this is not an option consumer goods firms have. Consequently, two conditions are required in order to reach zero-emissions[1]: the entirety of the electricity provided by the energy firm must be generated by green plants, and sector 1 firms must entirely electrify. For their part, the CO_2 emissions form the input for the DSK's climate component, which was initially simulated using the Climate Rapid Overview And Decision Support (C-ROADS) model, presented in [28]; however, for this work the climate model used is the Finite Amplitude Impulse Response (FAIR) model, for reasons detailed in the following section. The increased temperature resulting from the CO_2 emissions output by the DSK's sector 1 and electricity sectors does not cause any damages to the non-climate component of the DSK, unlike in the version of the DSK presented in [16]. It is partly as a consequence of this that the modelled temperature increase itself will be one of the outcomes used to judge the relative effectiveness of different policies. As noted in chapter 1, this is very much distinct from the method of policy evaluation used in models like DICE, in which it is assumed the damage inflicted on the economy's productivity in the model by rising temperatures is sufficient that the GDP-derived welfare is the only indicator needed to judge the effectiveness of different policies.

[1] The DSK model, in its current implementation, features neither negative emissions technologies nor opportunities for land-use change which might increase carbon sequestration. The latter point results from the fact that land use is not currently simulated by the DSK model. Therefore, it makes little sense to refer to net-zero emissions when discussing the DSK model; the term is thus eschewed in favour of 'zero-emissions'.

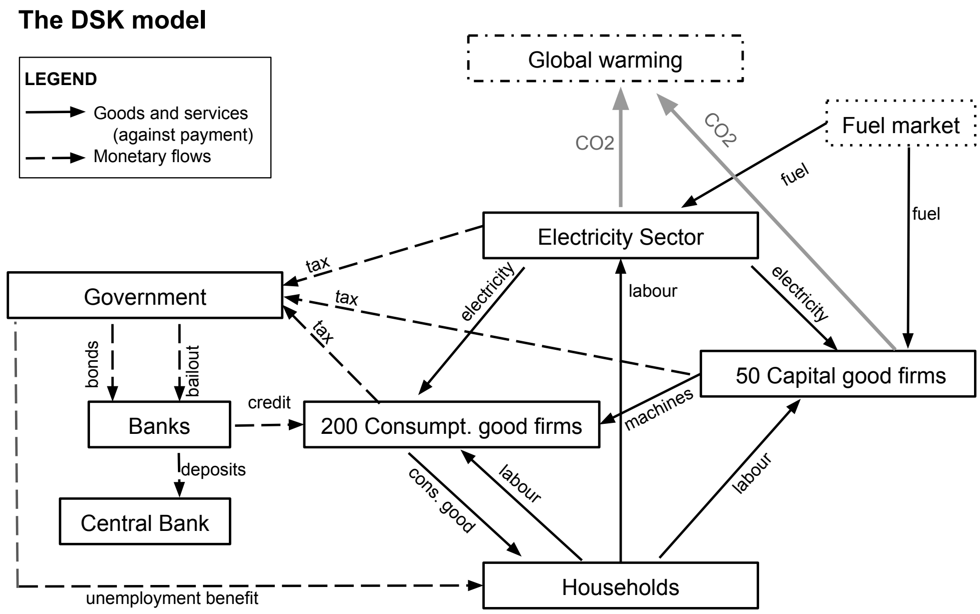

Figure 2.1: Schematic representation of the Dystopian Schumpeter meeting Keynes (DSK) model. Figure adapted from [27].

In addition to temperature change, unemployment will be used as the primary economic indicator of the relative merits of different policies[2]. As can be seen in fig. (2.1), all employment in the DSK model is concentrated in capital and consumer goods firms and the energy sector. Each timestep, the labour demand of each firm in the industrial sector is computed, as is the labour demand of the electricity firm. When the costs incurred on firms by a carbon tax are high, as will be the case in many of the simulations discussed in chapter 4, the aggregate labour demand will drop as a result of the decrease in labour demand in the individual firms. It should be noted that the labour demand of the electricity firm is also dependent on the firm's activities. Intuitively, the more green (or brown) plants being built in a given timestep, the higher the labour demand of the electricity firm. Additionally, if the number of green plants is being expanded very rapidly, as can be the case when there are far too few green plants to meet the demand for electricity and there is a high carbon tax, the price of each new green plant is higher than the last, rising according to

$$C_g(n) = C_g(1)\left(1 + \frac{1}{N_{lim}}\max\left\{1, n - N_{lim}\right\}\right), \tag{2.1}$$

where $C_g(n)$ is the cost of the n^{th} green plant built in a given timestep and N_{lim} the threshold beyond which this additional cost is applied. In this work N_{lim} is 20% of the existing green plant stock. Since all the costs associated with green plants in the DSK model are assumed to be labour costs, the building of each new plant beyond this threshold will require more labour than the last. Thus, in years when there is a massive expansion in the number of green plants being built, labour demand will increase steeply. It should be noted that the labour demand associated with new green plants is split evenly between the timestep in which the green plant is built and the green plant's lifetime, the latter representing maintenance costs. This split was a change made to the DSK model for this project: initially all green plant costs were associate with plant building[3]. Once labour demand in a given timestep has been met, the remainder of the working-age population is paid unemployment benefits by the government, funded through a combination of taxes and the issuing of bonds.

[2] Unemployment is favoured over the GDP as an indicator in this work for two main reasons. First, there is some controversy surrounding the use of GDP as a well-being indicator [29]. Second, polling in the USA and UK indicate that more people than not believe that more focus should be put on environmental protection, even at the expense of economic growth [30, 31]. There is little evidence, however, that this holds for unemployment. Nonetheless, the effects of the policies tested on GDP will be shown in much of the results section, so the reader may come to their own conclusions.

[3] The 50-50 split between building and maintenance was reached by taking data provided by the International Renewable Energy Agency (IRENA) and a report for the Canadian Hydropower Association on the jobs created by solar [32], wind [33, 34] and hydropower plants [35], and weighting them by the number of people currently employed for each form of electricity generation [36].

2.2 Initial Conditions and the Choice of Climate Model

As stated in chapter 1, the research questions this report addresses necessarily involve simulating policy-making strategies over a range of different ECSs, as the true ECS is still largely uncertain. One problem that presents itself when using a simple climate model, such as C-ROADS, is how to justify having the same conditions at the start of the climate and policy simulations - the year 2020 - for a range of different ECSs. In order to lay the groundwork for this discussion, though, a note will first be made on the way that warming is referred to in this report, and its relation to global mean surface temperature (GMST).

2.2.1 A note on baselines

Of central importance to any discussion involving the limits within which signatories of the Paris Climate accords have agreed to stabilise GMST is the baseline that is used [37]. Consequently, a note is made here that the 'temperature change' referred to in this work is always the temperature anomaly in degrees Kelvin (or Celsius), relative to the 1850-1900 GMST. This baseline corresponds to the *pre-industrial baseline* used by the World Meteorological Organization (WMO) and IPCC [38, 37].

2.2.2 ECS and TCR

According to the WMO's State of the Global Climate 2020 report, in 2020 GMST had increased by approximately 1.2°K, relative to the pre-industrial baseline. Consequently, a GMST anomaly of 1.2°K is the desired temperature initial condition for all policy simulations. However, this presents a problem in terms of the logical consistency of the model - if the ECS is the only measure of climate sensitivity that can be specified for in a given climate model, then that model, given the same dataset of historical CO_2e emissions, will yield different degrees of warming by 2020 for different ECSs. This is demonstrated in fig. (2.2(a)), which shows the modelled evolution of the temperature anomaly for all different ECSs resulting from historical CO2e emissions (including land use change). The data used for these simulations was retrieved from Our World in Data, and is originally from [39]. For an explanation of the irregularly-spaced ECS values simulated, see chapter 3.

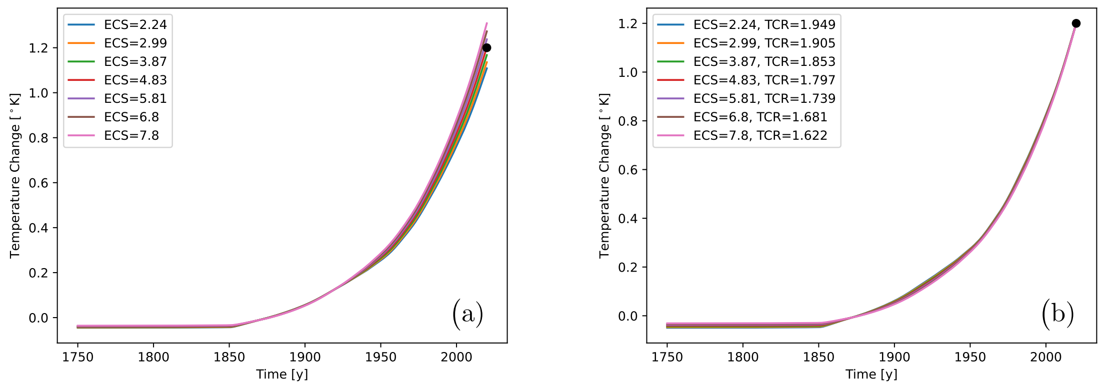

Figure 2.2: Simulated temperature change, relative to the 1850-1900 mean, in the FAIR model, resulting from historical CO2e emissions, for 7 different ECSs. Panel (a) shows the simulated result for a fixed TCR of 1.8°K/doubling CO_2, while (b) shows the modelled temperatures when TCR is adapted to each ECS to ensure that the 2020 temperature anomaly is 1.2°K. See chapter 3 for an explanation of the irregularly-spaced ECS values. The CO2e emissions data includes land use change and was retrieved from Our World in Data, originally from [39].

In order to make the climate model initialisation more logically consistent, the decision was made to use a climate model in which the Transient Climate Response (TCR) can be specified, in addition to, and independently of, the ECS. The TCR is defined as the GMST response, relative to the pre-industrial value, to a scenario in which CO_2 concentration increases from preindustrial levels by 1% every year, at the time of doubling (70 years after the start of the increase) [20]. As such, the TCR is a measure of immediate, as opposed to equilibriated, climate sensitivity, thus being also a measure of the speed of the climate's response to increased CO_2 concentrations. It follows, then, that in a model in which both the ECS and the TCR can be independently specified, there is a value of the TCR for every ECS that can assure that the temperature anomaly in 2020 is 1.2°K. The model used to this end was the Finite Amplitude Impulse Response (FAIR) climate model, which is an adapted version of a simple climate model developed for the IPCC's 5th Assessment Report, tuned to reproduce the behaviour of higher complexity Earth System Models [40].

2.2.3 The FAIR model

The FAIR model consists of a simplifed carbon cycle model, coupled to 2 temperature reservoirs by a simple radiative forcing equation, with the sum of the two temperature reservoirs yielding a value for the modelled GMST [40]. The carbon cycle consists of 4 carbon reservoirs, each with a carbon concentration anomaly R_i, the evolution of which is governed by

$$\frac{\mathrm{d}R_i}{\mathrm{d}t} = a_i E - \frac{R_i}{\alpha \tau_i}; i = 1 - 4 \tag{2.2}$$

where E represents the yearly carbon emissions, a_i the coefficient which determines the fraction of emissions taken up by the i^{th} reservoir, τ_i a timescale factor and α an additional coefficient which varies slightly each timestep and parametrises the 100-year integrated impulse response function[4]. The sum over the carbon reservoirs yields the total atmospheric carbon concentration, which in turn determines the radiative forcing:

$$F = F_{\text{ext\_factor}} \frac{F_{2\times}}{\ln(2)} \ln\left(\frac{C}{C_0}\right); C = C_0 + \sum_i R_i. \tag{2.3}$$

Note that F_{ext_factor} represents a non-CO_2 forcing factor[5], while $F_{2\times}$ is the forcing due to a doubling in CO_2. Finally, the evolution of the two temperature reservoirs is determined by

$$\frac{\mathrm{d}T_j}{\mathrm{d}t} = \frac{q_j F - T_j}{d_j}; j = 1, 2, \tag{2.4}$$

where d_j is the timescale associated with each temperature reservoir. The values for all parameters, excluding q_1 and q_2, are detailed in table (2.1). The reason for these three parameters being omitted from the table is that their values determine the ECS and TCR, through

$$q_1 = \frac{1}{F_{2\times}} \frac{A_2 ECS - TCR}{A_2 - A_1} \tag{2.5}$$

and

$$q_2 = \frac{ECS}{F_{2\times}} - q_1, \tag{2.6}$$

where A_1 and A_2 have been defined, for the sake of notational brevity, as

$$A_j = q_j \left(1 - \frac{d_j}{70}\left(1 - \exp\left[-\frac{70}{d_j}\right]\right)\right); j = 1, 2.$$

The relations in eqs. (2.5) and (2.6) were obtained by reformulating eqs. (4) and (5) in [40].

To find the TCR for each ECS, such that the 2020 temperature anomaly resulting from historical CO_2e emissions was constrained by the desired value, the FAIR model was run for 100 different ECSs between 1 and 8°K/doubling CO_2. For each of these values, the TCR was found that would minimise the difference in simulated 2020 temperature anomaly from the desired value of 1.2°K using `scipy`'s `optimize.minimize()` function. As the required TCR's relationship to the ECS turned out to be virtually indistinguishable from a linear function, a linear fit was found using a least-squares method, so that the TCR could be found for any ECS in that range. The resulting modelled temperature time series, with the TCR adapted to each ECS, is shown in fig. (2.2(b)); as can be seen, the temperature is successfully constrained by the desired 2020 value under every ECS. An important consequence of the choice to adapt the TCR according to the modelled ECS is that in higher ECS cases, GMST will take longer to stabilise once CO_2 emissions have reached zero, as will be seen in chapter 4.

[4]The 100-year integrated impulse response function represents the 100-year average airborne fraction of a pulse of CO_2; the authors of the FAIR model noted that this better reflects the impact of CO_2 emissions than the airborne fraction at any particular moment in time [40]

[5]In the FAIR model as presented in [40], the non-CO_2 forcing parameter was an additive constant, rather than a factor. However, as there is currently no source of non-CO_2 forcing outputted by the DSK model, and it seemed unwise to assume a constant value for the non-CO_2 forcing, the simplfying assumption is made that non-CO_2 forcing is proportional to the CO_2 forcing.

Parameter	Value - FAIR	Guiding analogues
a_0	0.2173	Geological re-absorption
a_1	0.2240	Deep ocean invasion/equilibration
a_2	0.2824	Biospheric uptake/ocean thermocline invasion
a_3	0.2763	Rapid biospheric uptake/ocean mixed-layer invasion
τ_0 (year)	1×10^6	Geological re-absorption
τ_1 (year)	394.4	Deep ocean invasion/equilibration
τ_2 (year)	36.54	Biospheric uptake/ocean thermocline invasion
τ_3 (year)	4.304	Rapid biospheric uptake/ocean mixed-layer invasion
q_1 (KW	-	Thermal equilibration of deep ocean
q_2 (KW m^2)	-	Thermal adjustment of upper ocean
d_1 (year)	239.0	Thermal equilibration of deep ocean
d_2 (year)	4.1	Thermal adjustment of upper ocean
$F_{2\times}(Wm^{-2})$	3.74	Forcing due to CO_2 doubling

Table 2.1: Parameters used in the FAIR model, as implemented in this work. The values for q_1 and q_2 are left blank in this table as these are used as a proxy for the desired ECS and TCR, and, as such, are different depending on the different values the policies are tested under. It should be noted that the guiding analogues are only rough interpretations, as the FAIR model is an empirical fit tuned to accurately represent temperature responses, and does not attempt to explicitly model the specific behaviours of each carbon and temperature reservoir. Table adapted from [40].

Chapter 3

Modelling and Evaluating Policy

In addressing the two main research questions outlined in chapter 1, a key addition to the DSK model was a form of policymaker whose understanding of the ECS, represented through the ECS PDF, updated as the model ran. This was achieved through a combination of Bayesian learning and simulated research on the ECS, which constrains the learning process, guiding it towards the simulated ECS.

3.1 Bayesian Learning, Realistic Noise and Research Functions

Bayesian inference, the primary method through which adaptive policymakers in this work update their understanding of the ECS, is based on the repeated application of Bayes' theorem,

$$p(\theta|y) = \frac{p(\theta)p(y|\theta)}{p(y)}, \tag{3.1}$$

on a probability distribution [41]. In this notation, θ corresponds to a hypothesis - in this context a possible value of the ECS - while y corresponds to a measurement - in this context, a temperature measurment. Thus, the ECS PDF, noted $p(ECS)$, is updated each timestep through a process of Bayesian inference based on the GMST measured in that timestep. Noting that the PDF is discretised into values labelled ECS_i, evenly separated in steps of $dECS = 0.005°\text{K/doubling CO}_2$, the probability density for each ECS_i is updated each timestep, as

$$p(ECS_i)_{t+1} = \frac{p(ECS_i)_t p(t|ECS_i)}{Z}. \tag{3.2}$$

Note that $p(y)$ from eq. (3.1) does not need to be explicitly calculated, but can rather be taken as a normalisation constant, to ensure that the definite integral of $p(ECS)$ over the whole domain is equal to 1. Bearing in mind the discretised nature of the PDF, the normalisation constant, denoted Z in eq. (3.2), is simply

$$Z = \sum_i \tilde{p}(ECS_i)_{t+1} = \sum_i p(ECS_i)_t p(t|ECS_i).$$

Measurement error is assumed to be negligible. Each timestep, for each potential ECS considered, the expected temperature is calculated using the FAIR model, using the history of CO_2 concentrations [1]. The policymaker is assumed to know the true TCR. Consequently, for $p(t|ECS_i)$, the probability of measuring a temperature t when the ECS is ECS_i, to not take the form of a Dirac delta function, stochastic noise had to be added to the temperature signal. This noise can be seen as representing atmosphere-ocean interactions, such as the El Niño Southern Oscillation.

3.1.1 Introducing noise to the FAIR model

The decision was made to make the noise broadly similar to inter-annual variations in the GMST measured in the Earth system, in order to simulate potential future Bayesian learning processes as faithfully as possible. To this end, it would be desirable to tune the frequency spectrum of the noise produced by the adapted FAIR model to the real-world GMST noise frequency spectrum. This would ensure that the variability in the temperature signal is seen at realistic scales over realistic timescales. However,

[1] In fact, the code written for this work followed a slightly different method, due to personal error. For more information see section 3.1.3

given the number of assumptions that would have been involved in isolating noise from the idealised true warming signal in real world datasets - if this were possible without making any prior assumptions, the uncertainty surrounding the ECS would likely be negligible - an Earth System model set to simulate surface temperature under constant CO_2 concentrations was used instead. Specifically, the dataset used for the tuning was the output of Community Earth System Model Version 2 (CESM2) under pre-industrial control settings, which keep global CO_2 concentrations at pre-industrial levels [42].

The introduction of noise to the FAIR model was achieved through the addition of two stochastic terms to eqs. (2.4), the equations which govern temperature in the model. Specifically, normally-distributed random noise terms were added to the differential equation governing the second temperature reservoir - that is, the one with the 4.1-year timescale - and the equation which sums over the two temperature reservoirs to yield the GMST analogue outputted by the FAIR model. The reason for this combination is that the latter stochastic term alone would not affect the evolution of the temperature at all: the system has no memory of this term, and the resulting white noise would have a frequency spectrum with uniform intensity [43]. The former stochastic term, however, does impact the temperature's evolution - consequently, the combination of the two terms, weighted appropriately, can be tuned to offer a reasonable approximation of CESM2's noise spectrum. With the differential equations discretised according to the forward-Euler method and Euler-Maruyama method [44] - for the non-stochastic and stochastic differential equations, respectively - the adapted equations read

$$\begin{aligned} T_{1,t+1} &= \frac{dt}{d_1}\left(q_1 F - T_{1,t}\right) \\ T_{2,t+1} &= \frac{dt}{d_1}\left(q_1 F - T_{2,t}\right) + \sigma_a dW_t \\ T &= T_1 + T_2 + \sigma_b X, \end{aligned} \tag{3.3}$$

where W_t is the Wiener process [44], and X a random variable following a standard normal distribution. The values $\sigma_a = 0.078$ and $\sigma_b = 0.024$ led to a noise spectrum that was relatively similar to the CESM2 spectrum, as can be seen in fig. (3.1).

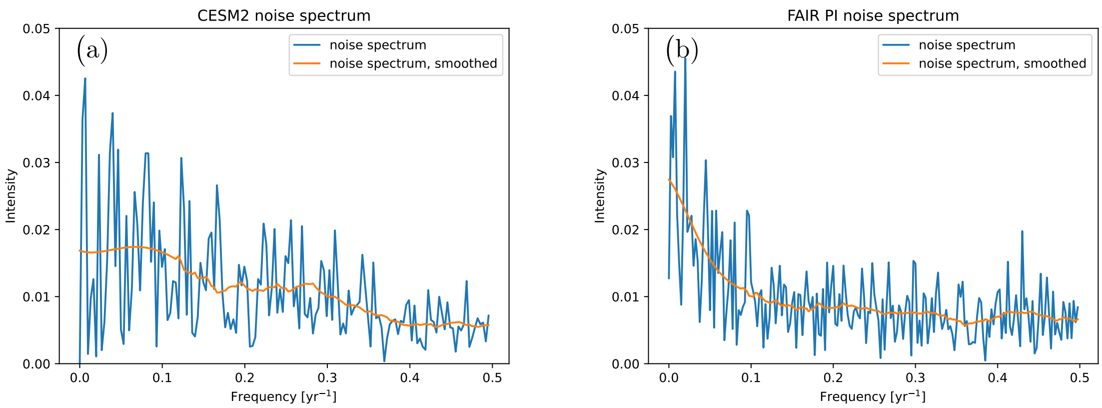

Figure 3.1: Fast Fourier Transform of a 400-year time series of temperature with constant, pre-industrial atmospheric carbon concentrations outputted by (a) CESM2 and (b) the FAIR model with the added stochastic terms. The orange lines in each panel shows a smoothed noise spectrum, obtained by applying a Savitzky–Golay filter, so that the two spectra can be more easily compared. It should be noted that, due to the finite nature of the time series and stochastic nature of the program, the noise spectrum outputted by the adapted FAIR model does not always look exactly as it does in panel (b). The variation between different simulations is greatest for lower frequencies, particularly as $\nu = 1/400yr^{-1}$ is approached. Nonetheless, the spectrum shown in panel (b) is a broadly representative example.

3.1.2 The learning process

The FAIR model is run each timestep for every potential ECS considered, ranging from 0.1°K/doubling CO_2 to 10°K/doubling CO_2 with a gridsize of 0.005° K/doubling CO_2. As stated, the policymaker is assumed to know the true TCR. However, it is not assumed that the policymaker is aware of the exact structure of the FAIR model's noise terms: the policymaker instead assumes a white noise spectrum, in which a normally-distributed random term is added to the temperature each year, with a standard deviation of 0.1°K, chosen as it is the standard deviation of the temperature in the CESM2 time series.

$p(t|ECS_i)$ is thus a Gaussian distribution with standard deviation of 0.1°K, centred around the temperature predicted by the strictly deterministic FAIR model for a given ECS. The initial PDF is a lognormal fit of the ECS PDF detailed in Sherwood et al.'s 2020 estimate [22].

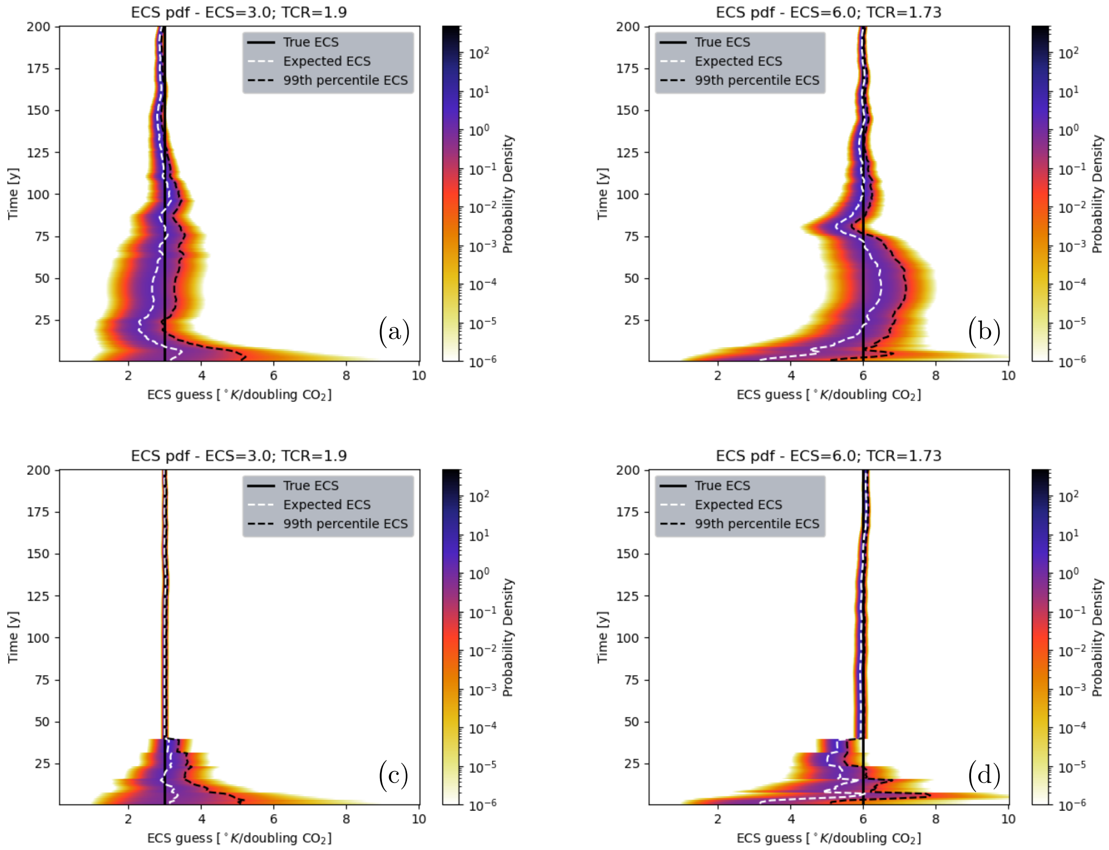

Figure 3.2: Time evolution of the ECS PDF using different methods and under different true ECSs. In each panel, the probability density is plotted as a heatmap with a logarithmic colour scale, shown on the right-hand side of each panel. Panels (a) and (c) show the evolution of the ECS PDF in the situation where the model's true ECS is 3°K/doubling CO_2, with (b) and (d) showing the evolution under a true ECS of 6°K/doubling CO_2. (a) and (b) show the evolution of the PDF when Bayesian inference is the only method through which new information about the ECS is learned and incorporated into the PDF; (c) and (d) show the evolution when the Bayesian inference process is supplemented with the research functions described in section 3.1.2. In all cases, the climate simulation starts from 2020 conditions, and the carbon concentration scenario modelled is one in which atmospheric CO_2 concentration increases each year by 0.5%: according to the data available at [45], this approximately corresponds to a continuation of the trend of the past 40 years. In each panel, the thick black line shows the modelled true ECS, with the white and black dashed lines, respectively, showing the expectation and 99^{th} values of the ECS.

Panels (a) and (b) of fig. (3.2) show the evolution of the ECS PDF under ECSs of 3°K/doubling CO_2 and 3°K/doubling CO_2, respectively. In both cases, the atmospheric CO_2 concentration starts at 2020 levels, as does the warming, and increases by 0.5% per year. Even under this somewhat pessimistic CO_2 concentration trajectory - according to the NOAA's climate globally averaged CO_2 trends [45], this trajectory is approximately equivalent to a continuation of the trend of the last 40 years - it takes 50 years for the expected ECS to consistently stay within 0.5°K/doubling CO_2 of the true value in the lower ECS case. In the higher ECS case this takes almost a century. Considering the width of the PDF, it takes a century for the 99^{th} percentile to reach similar levels of agreement with the true value in the lower ECS case, and a similar amount of time in the higher ECS case.

Given that this project will compare the outcomes of policymakers with differing levels of risk-aversion who adapt to changing ECS knowledge, panels (a) and (b) of fig. (3.2) show that these policymakers will likely have significantly different policies for a period of around a century. This timeframe is judged to be too long for the purposes of this project. The reasons for this are twofold. First, the coming decades are the most crucial in terms of abating CO_2 emissions [6], so it makes more sense that divergences in different policy approaches be most significant in the near future, as opposed to over an entire century. Second, the assumption that the policymaker would not have access to any research beyond the latest measurements of GMST is difficult to justify. The very existence of the IPCC's Working Group 1's reports, themselves reflective of a vast body of climate research including research - including ECS estimations as have already been cited in this report [21] - demonstrate the unrealistic nature of this assumption. Consequently, the choice has been made to constrain the evolution of the ECS PDF with the publication of 5 'researched' PDFs, over 40 years, which are incorporated into the policymaker's pdf as additional Bayesian learning steps.

The researched PDFs, as with the initial ECS PDF, are defined by a lognormal function, which can be defined in terms of the parameters μ and σ as [46]

$$\rho(x;\mu,\sigma) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln(x)-\mu)^2}{2\sigma^2}\right). \tag{3.4}$$

Note that μ and σ are not the mean and standard deviation, respectively, of the lognormal distribution: these are given by [46]

$$\begin{aligned} \mu_x &= e^{\mu+\frac{1}{2}\sigma^2}, \\ \sigma_x &= e^{\mu+\frac{1}{2}\sigma^2}\sqrt{e^{\sigma^2}-1}. \end{aligned} \tag{3.5}$$

The initial PDF used is simply the initial PDF for the Bayesian learning. Each subsequent researched PDF linearly approaches the true ECS, such that the peak of the distribution reaches the true value in the last PDF, while the standard deviation of the distribution reaches a value of one step in ECS ($dECS$) in the final PDF. Due to the narrowness of the distribution once the final researched PDF has been incorporated, subsequent Bayesian steps have little effect on the PDF after this point. The four panels of fig. (3.2) show the evolution of the ECS PDF, with and without the research functions. The research functions have the desired effect of constraining the learning process so that it is effectively complete after 40 years, by 2060.

3.1.3 Difference between the theory and the implementation

Due to personal oversight, there is a discrepancy between the way the Bayesian learning process is described in previous sections and the way it was implemented in the code used for this project. The difference lies in the calculation of the expected temperature, for each ECS considered, each timestep. The method chosen was to update these expected temperatures each timestep from the expected temperature of the previous timestep, for each ECS. However, in the code, a new expected temperature is developed each timestep for each ECS *from the actual temperature measured in the last timestep.* Qualitatively, this is likely to lead to slower learning, as the noisy true temperature signal affects the expected temperature in the following timestep, leading to a less consistent evolution in the expected temperatures, stymying the learning process. However, due to the addition of the research functions, this is unlikely to significantly affect the evolution of the expected and 99^{th} percentile values of the ECS - and, by extension, the strategies of the adaptive policymakers. This is demonstrated in fig. (3.3), which can be used to compare the evolution of the ECS PDF under two ECSs with correctly and wrongly implemented code. Nonetheless, this is a mistake and as such should be corrected if the work presented in this report is to be taken further.

3.2 Policies and Policymakers

Recalling the aims laid out in the introduction, it should be noted that this work's research questions make no assumptions about the type of policy that should be followed. Furthermore, there is no reason to assume that the answers to these questions will be the same for all types of policies. Consequently, different policymakers are compared while applying different types of policy.

Two types of policy are focused on. The first is a carbon tax, with no other policy; in the second, a carbon tax funds the building of green plants by the government. In the case of the second policy, 50% of the carbon tax revenue funds a green plant-building scheme, with the other 50% left free for other purposes not considered explicitly in the model - this might include providing additional funds for public transport, schools or hospitals, for example. The carbon tax is set at the same level for both policies: consequently, we expect to see CO_2 emissions abatement to be greater for the second policy than the first, all other factors being held constant. These policies will henceforth be labeled P1 and P2, respectively. For both P1 and P2, the carbon tax is ramped up over the first 5 years of the model run to the level the policymaker wishes it to be.

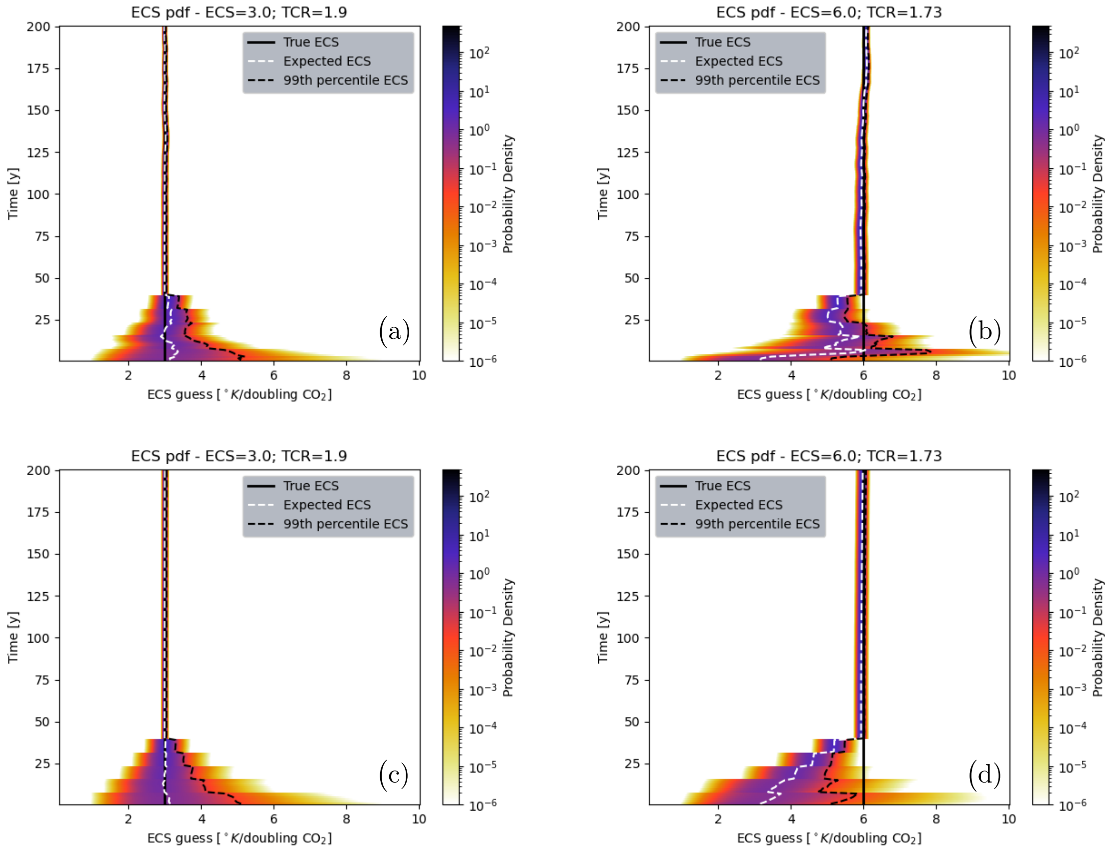

Figure 3.3: Evolution of the ECS PDF when Bayesian inference is supplemented with researched PDFs, under two ECSs, shown in the left and right panels as in fig. (3.2). The top panels, (a) and (b), are identical to panels (c) and (d) in fig. (3.2), showing the evolution of the PDF when the Bayesian inference process is implemented as intended. Panels (c) and (d) of this figure show the evolution of the PDF with the mistaken method used in the adapted DSK model in this project. Qualitatively, there is little visible difference in the evolution of the expected and 99$^{\text{th}}$ percentile values of the ECS between the correctly and wrongly implemented code. It should be noted that, due to the stochastic nature of the temperature signal, the evolution of the PDF is different each time the code is run. Consequently, when comparing (b) and (d) it should *not* be concluded that the 99$^{\text{th}}$ percentile ECS value never overshoots the true ECS in the wrongly implemented code: as is seen in the figures in chapter 4, this is clearly not always true.

For each policy, four policymakers are considered:

1. Risk-neutral, non-adaptive;
2. Risk-averse, non-adaptive;
3. Risk-neutral, adaptive;
4. Risk-averse, adaptive.

Non-adaptive policymakers impose a constant carbon tax - with the exception of the five-year ramp-up - over the whole model run, based on their level of risk aversion. Adaptive policymakers, by contrast, change the carbon tax level according to their evolving knowledge of the ECS. The difference between risk-neutral and risk-averse policymakers is implemented in the model through the percieved, or virtual, ECS, labeled VECS. This is the ECS, identified by some property of the ECS PDF, that the policymaker takes as their estimate of the true ECS. In other words, the policymaker choses to focus on a position on the ECS PDF, determined by their level of risk-aversion, and sets their carbon tax according to that ECS. Risk-neutral policymakers use the expectation value of the ECS, with risk-averse policymakers focusing on the 99$^{\text{th}}$ percentile value. A non-adaptive, risk-averse policymaker will therefore set their carbon tax as if the ECS were the 99$^{\text{th}}$ percentile of the Sherwood et al. PDF [22] and maintain this level over the course of the model run.

The carbon tax scales with VECS according to

$$CT = cVECS^2, \tag{3.6}$$

where c is a coefficient that has been calibrated such that a non-adaptive, risk-averse policymaker using P1 is able to consistently keep warming by 2100 to under 2°K, if the true ECS is 2.99°K/doubling CO_2. The quadratic relationship has been chosen so that the warming that occurs under adaptive policymakers does not increase too drastically if the ECS is increased. Keeping the warming seen constant with respect to ECS for the adaptive policymakers was not used as a strict constraint, as the carbon tax would need to be so high that unemployment would regularly exceed 50%. For this reason it seemed unreasonable to maintain such a constraint, especially considering the fact that P1 only uses one of the mitigating policies available to the policymaker.

3.3 Evaluating Over Different ECSs

All policymakers - and all policies - must be evaluated under a range of different values of the ECS. This leads to difficulty when it comes to evaluating that policymaker and policy overall. While the outcomes of different policymakers following different policies will be shown in chapter 4 for several ECSs, another approach which will be used is to aggregate these outcomes over all ECSs tested. A risk-neutral way of aggregating the outcomes of a given policy is to take the probability-weighted average of the indicator of interest, I_k, over the N different ECS values, ECS_i tested:

$$\bar{I_k} = \sum_{i=1}^{N} I_{k,i} P(ECS_i). \tag{3.7}$$

Note that $P(ECS_i)$ here refers to the probability of the ECS falling between $(ECS_{i-1} + ECS_i)/2$ and $(ECS_i + ECS_{i+1})/2$. If each ECS_i considered is separated by a step in ECS, labelled $dECS'$, $P(ECS_i)$ can be approximately computed by numerically integrating the PDF as a Riemann sum with the more fine-grained step in ECS, $dECS$,

$$P(ECS_i) = \sum_{ECS_j = ECS_{i-1/2}}^{ECS_{i+1/2}} ECS_j \cdot \rho(ECS_j) \cdot dECS,$$

where $ECS_{i\pm 1/2}$ is short-hand for $ECS_i \pm dECS'/2$. In theory, then, $\bar{I_k}$ should be equivalent to the expectation value of I_k, according to the estimated ECS PDF in the Sherwood et al. report [22].

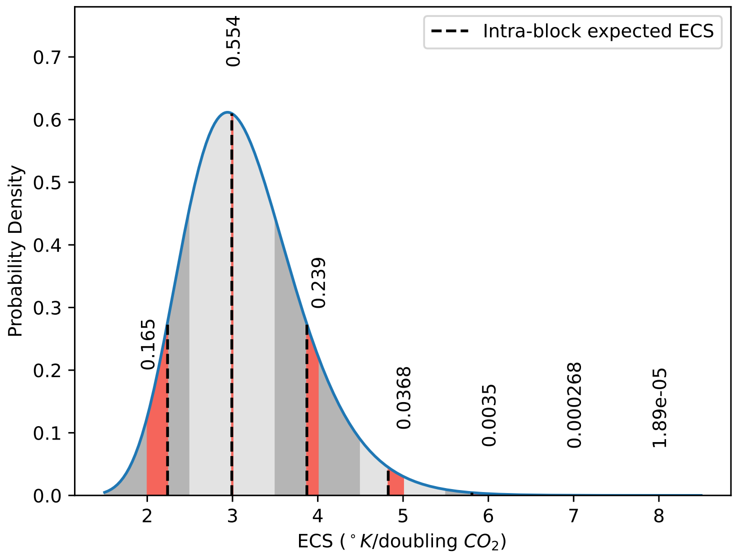

Figure 3.4: The lognormal approximation to the ECS PDF found in [22], split into blocks of width $dECS' = 1°K$/doubling CO_2. The dotted, vertical lines correspond to the expected ECS within each block; the discrepancy between this and the centre of the block is shaded red for emphasis. The probability of each block is written above the PDF.

However, one problem with the method for aggregating indicators over different ECS values described in eq. (3.7) is that the probability density is not evenly distributed within the blocks of width $dECS'$.

Figure (3.4) demonstrates the differences between the expected ECS and block centre in each block. While the discrepancy between the two appears negligible near the peak of the distribution, effects become more significant further from the peak. Accordingly, the DSK is instead run with ECS values corresponding to the intra-block ECS expectation values, so that the binned distribution is not too poor an approximation of the full pdf. The probability-weighted sum over different ECSs should then be a more accurate estimated expectation value.

As has been noted, this is a risk-neutral way of evaluating the effectiveness of different policymakers and policies: a risk-averse form of aggregation would involve weighting the tail of the ECS distribution, and using this modified distribution to weight the different simulations. Risk-averse methods of evaluating policy are not used in this work, due to the arbitrary nature of the choice of weighting function needed to give greater weight to higher ECSs. Nonetheless, it is important to acknowledge the fact that by weighting policies in the risk-neutral way, strictly according to the estimated probability of each ECS, a certain point of view will be implicit in the conclusions this report reaches. No matter how grave, the results of a given policy under a very high ECS - for example 7.8°K/doubling CO_2, the highest ECS simulated - will be given no more weight than the 0.00189% chance that the ECS lies between 7.5 and 8.5°K/doubling CO_2. That this method will be applied, regardless of how socially and/or environmentally destabilising the outcome of the policy under this ECS is, is a choice: while justifiable, the reader may not necessarily agree.

Chapter 4

Results

4.1 Baseline: No-Policy

Before considering the effects of different policies and attempting to judge their relative effectiveness, the behaviour of the DSK model when no climate policy is enacted is first considered. This should enable more effective judgement of what good and bad outcomes in the domain of unemployment are, within the context of the model. All simulations presented in this chapter have been run for 50 Monte Carlo realisations.

Panels (a) and (b) in fig. (4.1) show how the temperature change and unemployment rate, respectively, evolve over the course of the model's run from 2000 to 2120; recall that the climate module, and climate policy when implemented, is only run from 2020 onwards. Figure (4.1) shows the results in the case that ECS=2.99°K/doubling CO_2. In addition to temperature change and unemployment rate, fig. (4.1) also shows the evolution of the share of electricity generated by green plants, the fraction of the capital goods sector (sector 1) that has electrified, the GDP and the CO_2 emissions themselves in panels (c), (d), (e) and (f), respectively. The plots of the electricity mix, electification in sector 1 and GDP are shown to help gauge the extent to which emissions abatement is due to a green transition in the model's energy system, as opposed to temporary economic contraction. Finally, the plot of CO_2 is included as it directly shows the impacts of the different policies and policy strategies on emissions abatement, without being influenced by the effects of the different ECS and TCR values on the climate module.

Observing fig. (4.1 (c)), it is apparent that without any climate policy, the building and operating costs of green plants never become cost competitive enough with brown plants for any electricity to be supplied by green plants: without climate policy, there is no transition towards renewable energy in the DSK. This is in accordance with previous studies that use the DSK model [27]. This may partly be a result of the fact that the model is initialised with 100% of electricity being supplied by brown plants. An additional factor could be the fact that there is no limit to the amount of fuel that can be burnt in the DSK model. It is also apparent from panels (d), (e) and (f) that this lack of any shift from brown to green electricity generation, in conjunction with the negligeable upwards trend in sector 1 electrification and the near-exponential increase in production, ensure that CO_2 emissions do not peak prior to 2120. The consequence of this is that temperature rises at an ever increasing rate over the model run, surpassing 5°C of warming by 2100 in approximately 90% of realisations, with an expected warming of approximately 7°C in 2120.

It is important to note that the unemployment rate rises over the course of the model run, with the mean value increasing from an average of approximately 4% in the first decade of the model's run to 7-8% in the last 2 decades. When the unemployment will be used to judge the effectiveness of each policymaker under all ECS conditions in sections (4.2.3), (4.3.2) and (**??**), the difference in unemployment, relative to the baseline case, will be considered. This is partly to avoid penalising a policy for having non-zero unemployment even in the case where the unemployment is lower than in the baseline case, and partly to counter the effect of the increasing baseline unemployment rate.

4.2 P1: Carbon Tax Only

For the sake of brevity, particularly to allow us to consider the results of policy across the full range of possible climate sensitivities simulated, the in-depth discussion of the time series under different policy-

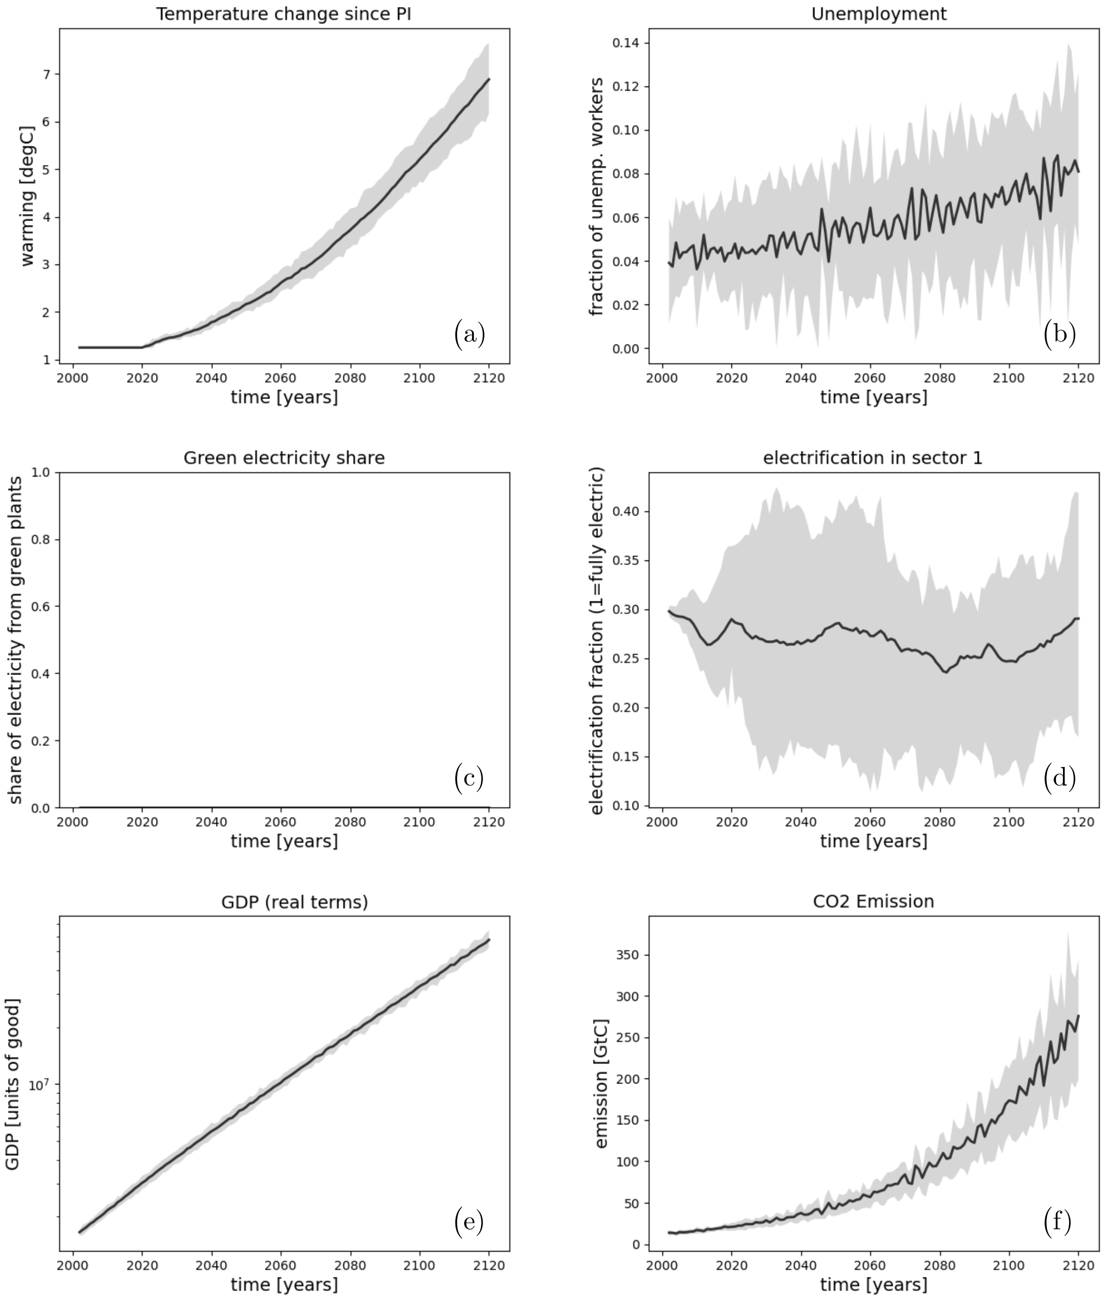

Figure 4.1: Evolution of the climate and economy from 2000 to 2120, in the case of no climate policy, if ECS=2.99°K/doubling CO_2. The model has been run for 50 realisations; the shaded region in each plot corresponds to the values bounded by the 10^{th} and 90^{th} percentiles, while the dark line shows the mean time series. Shown are (a) the global mean surface temperature change with respect to pre-industrial levels, modelled from 2020; (b) the unemployment rate; (c) the share of electricity produced by green plants; (d) the fraction of the capital goods sector which has electrified; (e) the GDP, adjusted for inflation, and (f) annual CO_2 emissions. Note that panel (e) is plotted with a logarithmically-spaced y-axis, so the near-linear trend corresponds to an approximate growth rate of 3% per year, as in [27].

makers will only be undertaken for P1.

4.2.1 Lower ECS case

Figure (4.2) shows the evolution of the carbon tax, in the DSK's unit of 'goods', shorthand for the consmer good sector's generic good, for the four policymakers under an ECS of 2.99°K/ doubling CO_2. As outlined in chapter 3, the two non-adaptive policymakers, referred to as 'fixed' in the fig. (4.2)'s legend, have a constant carbon tax in real terms - corrected for inflation - after the first five years of active policy, in which the tax is ramped up to the final value. The ECS of 2.99°K/ doubling CO_2 is only slightly under the expected value of the PDF in Sherwood et al.'s estimate [22]. This is why the adaptive, risk-neutral policymaker's carbon tax remains very close to that imposed by their non-adaptive counterpart, reaching a final value in 2060 slightly below the fixed, risk-neutral tax.

Now considering the adaptive, risk-averse policymaker, the tax is initially high, only slightly under the level imposed by the non-adaptive, risk-averse policymaker, before decreasing to join the value of the taxes imposed by the risk-neutral policymakers. Recalling fig. (3.2), this is a result of the fact that as the ECS PDF collapses around the expectation value, the 99^{th} percentile ECS is brought closer to this, until the two values are virtually the same - the difference between them being of the order of 0.01°K/ doubling CO_2 - by 2060, the year in which the learning is effectively complete.

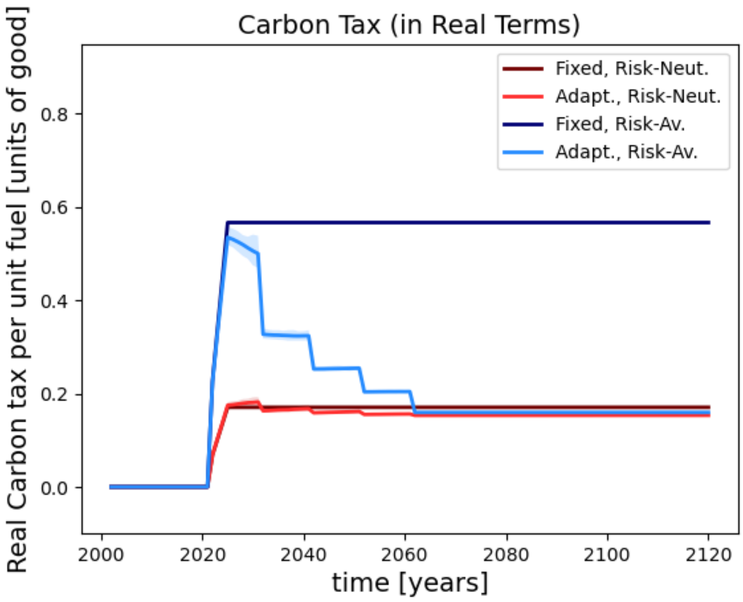

Figure 4.2: Evolution of the carbon tax under different policymakers under an ECS of 2.99°K/doubling CO_2, when the policy pursued is P1 (carbon tax only). The four types of policymaker investigated are shown, with red colours corresponding to risk-neutral policymakers and blues corresponding to their risk-averse counterparts. Lighter lines correspond to policymaking which adapts to new knowledge of the ECS, while darker lines correspond to policymakers who do not change their approach as time passes (the fixed case). As in fig. (4.1), the model has been run for each case for 50 realisations, with the shaded regions showing the values bounded by the 10^{th} and 90^{th} percentiles, and the solid lines showing the means.

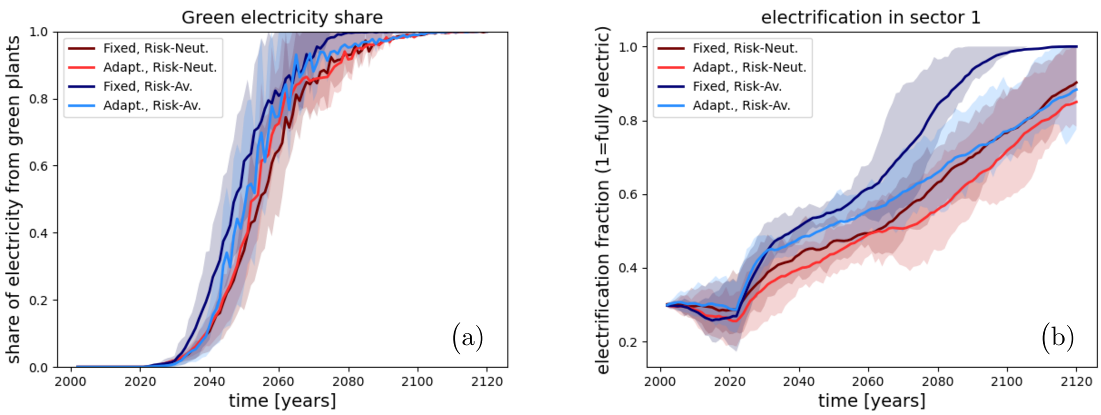

Figure 4.3: Evolution of (a) the percentage of electricity supplied by green plants and (b) the fraction of the capital good sector which has electrified, from 2020 to 2120 for different policymakers when P1 (carbon tax only) is pursued, under an ECS of 2.99°K/doubling CO_2. The colour scheme for the 4 policymakers shown is the same as in fig. (4.2).

Turning our attention to the results of the carbon taxes shown in fig. (4.2), fig. (4.3) shows the evolution of the percentage of the electricity supply generated by green plants alongside the fraction of the capital good sector which has electrified, in panels (a) and (b), respectively. Observing panel (a), it seems that, as a general rule, the policymakers with higher carbon taxes oversee a faster green transition in energy generation. The fixed, risk-averse policymaker is the first to see a full transition, ahead of other policymakers by approximately 10 years; the adaptive, risk-averse policymaker leads the risk-neutral policymakers by a few years at almost every stage, even in 2060-2080, when they have the same carbon tax. This is likely a result of the fact that an early, high carbon tax in the DSK model ensures that green plants become cost competitive with brown plants early, leading to a substantial increase in green plant building, which in turn leads to a more rapid decrease in the price of green plants from increased innovation. Consequently, there are more green plants by 2060 than under risk-neutral policymakers; furthermore, even when the carbon tax is decreased, green plants will have a lower relative cost when compared with brown plants than they would have if the carbon tax had not previously been so high.

An important caveat to this point is that the initial uptick in green plant building seen under the risk-averse policymakers is no earlier than under the risk-neutral policymakers. Observing fig. (4.3(a)) closely, all four policymakers see approximately the same green electricity share until 2030, depite their respective carbon taxes having reached the full value in 2025. This is partly due to the fact that, when initially imposed, the high carbon tax an effect on the energy sector's finances, meaning that there is less money to invest in building the green plants. Perhaps a more important dynamic, though is the effect on production, and hence fuel electricity demand, of the high initial carbon taxes, as will be seen when discussing the GDP.

Observing fig. (4.3(a)), the relationship between carbon tax and electrification is perhaps more straightforward. From the moment the climate module and policy routines in the DSK model start running, in the model year 2020, policymakers with higher carbon taxes invariably have a higher rate of electrification. This leads to a key difference among risk-averse policymakers, between the adaptive and non-adaptive cases. While, as we have noted, an initially high carbon tax may be enough to ensure slightly greater green plant adoption by the electricity firm even after it is reduced, the carbon tax must be maintained at a high level to ensure faster electrification. This is seen in the way that the electrification fraction of the adaptive, risk-averse policymaker joins that of the risk-neutral policymakers after 2060, when the carbon taxes they impose are very similar. Clearly, the rate electrification is highly dependent on size of the carbon tax, and it is not the case that a high initial carbon tax will give any degree of momentum to the electrification process. Considering the way in which the carbon tax incentivises electrification, this makes some sense: capital good firms' method of generating electricity is slightly less efficient than the electricity firm's brown plants. Consequently, even when the electricity firm has not yet undergone a full energy transition, a sufficiently high carbon tax will make electrification the cheaper option for sector 1 firms - the extent to which it is cheaper, and hence the rate of electrification, is directly tied to the severity of the carbon tax. It is only when 100% of electricity is generated by green plants that the electrification rate might accelerate without an increase in carbon tax; however, even here some carbon tax may be necessary to ensure that all firms electrify.

More generally, it should be noted that, with the current carbon tax calibration and when the carbon tax is the only climate policy implemented, only the fixed, risk-averse policymaker attains full electrification by 2120. Consequently, only this policymaker successfully decarbonises by the end of the model run.

Figure (4.4) shows the evolution of GDP under the four policymakers. The carbon taxes imposed by both risk-averse policymakers cause recessions, which end shortly after the end of the 5-year ramp-up period. The risk-neutral policymakers' carbon taxes cause at most short-lived recessions; in fact, the GDP resumes its growth before the ramp-up period is over. The recessions seen under the risk-averse policymakers are severe enough that the GDP only recovers to its 2020 levels approximately 10 years later. Notably, this coincides with the time that the green electricity share seen under the fixed, risk-averse policymaker outgrows that of the other policymakers. From this, it might be deduced that the four policymakers have similar green electricity shares during the 2020s because the recession depresses demand for electricity. Consequently, electricity demand is already met with the existing brown plant stock: despite the higher carbon taxes of the risk-averse policymakers, there is thus little rationale for a particularly rapid expansion of the green plant stock during this period.

The evolution of the CO_2 emissions under the four policymakers is shown in fig. (4.5). Confirming what was already noted when discussing fig. (4.3), the fixed, risk-averse policymaker is the only one consistently successful in reaching zero emissions by the end of the model run. The other policymakers each reach zero emissions in approximately 10% of realisations in 2120, but still have significant emissions in the the bulk of realisations. Also visible in fig. (4.5) is the effect that the recessions have on emissions. CO_2 emissions sharply decrease during the first few years of the carbon tax's implementation - especially strongly in the case of the risk-averse policymakers, whose larger and more protracted recessions are related to stronger decreases in energy use and hence emissions. The emissions increase following this

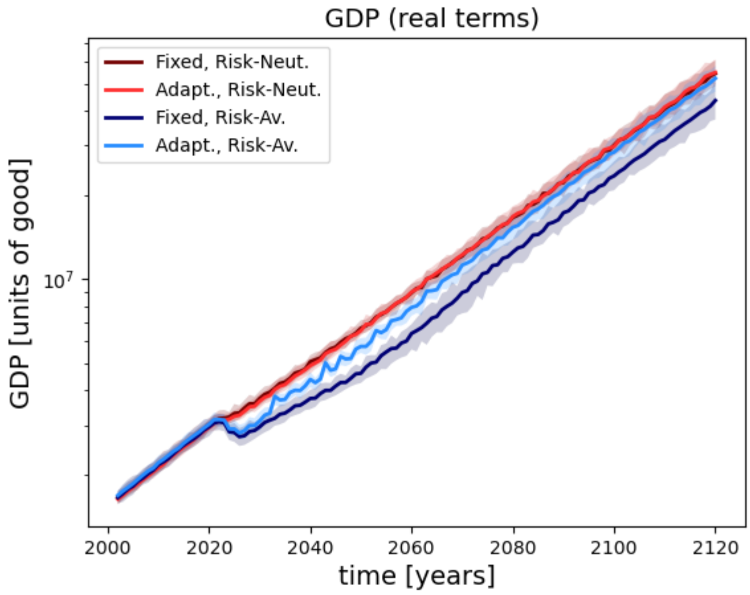

Figure 4.4: DSK model's GDP from 2020 to 2120 for different policymakers pursuing P1, under an ECS of 2.99°K/doubling CO_2. The colour scheme for the 4 policymakers shown is the same as in figs. (4.2) and (4.3).

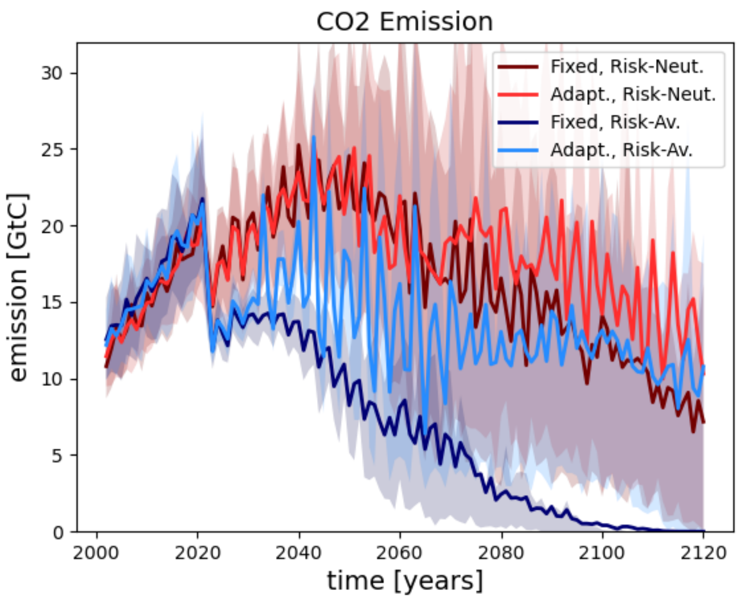

Figure 4.5: Modelled CO_2 emissions from 2020 to 2120 for different policymakers pursuing P1, when ECS=2.99°K/doubling CO_2. The colour scheme for the 4 policymakers shown is the same, as in figs. (4.2)-(4.4).

initial dip under all policymakers; however they quickly peak under the fixed, risk-averse policymaker due to the more advanced stage of green electricity generation and electrification transitions under that policymaker, whereas emissions continue to increase for longer under the other policymakers. The strong oscillations seen in fig. (4.5), most pronounced for the adaptive, risk-neutral policymaker, reflect the changes in energy use which follow the model's business cycles, which are just visible in fig. (4.4).

Figure (4.6) shows the evolution of both the temperature and unemployment under each policymaker, from which the performance of the different policymaking stragies - when ECS=2.99°K/doubling CO_2 and following P1 (the pure carbon tax) - may be judged. Considering first the temperature, in panel (a), we note that only the risk-averse, non-adaptive policymaker manages to stabilise the climate by the end of the model run, keeping warming visibly under 2°K in 90% of realisations. As expected, the difference between the adaptive and non-adaptive, risk-neutral, policymakers is minimal, with the expected warming

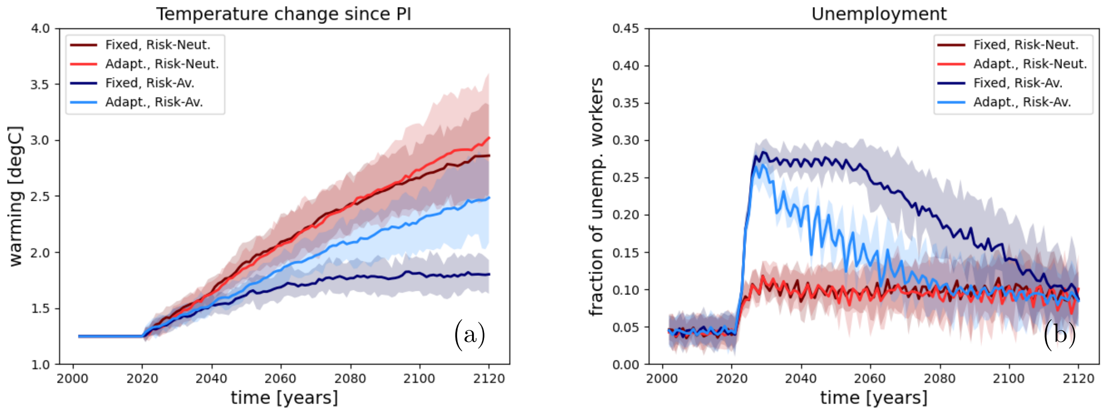

Figure 4.6: Evolution of (a) GMST change from pre-industrial levels and (b) unemployment rate from 2020 to 2120 for different policymakers implementing P1 when ECS=2.99°K/doubling CO_2.

of the two being virtually identical in 2100 and differing by approximately 0.1°K by 2120. Both see a fairly catastrophic temperature change with respect to pre-industrial of 3°K at the the end of the model run.

Among the risk-averse policymakers, there is a significant difference between the fixed and adaptive policymakers. As noted, only the non-adaptive policymaker consistently stabilises the climate, indeed meeting a key Paris Climate Accord goal, while the adaptive, policymaker sees approximately half a degree more warming. As noted this is largely due to the large difference in electrification between the two, due to the adaptive policymaker decreasing the carbon tax before the electricity generation transition and electrification are complete, although the fact that the GDP increases at a faster rate under them also plays a role.

It could be argued, upon observing the poor climate mitigation achieved by most modelled policymakers, that the carbon tax simply is not high enough in most cases. However, considering the unemployment, shown in panel (b), it becomes difficult to argue this point. Observing panel (b), it can immediately be remarked that the unemployment rises to extremely high levels, particularly under the risk-averse policymakers. On average, unemployment under these policymakers peaks at just under 30% of the working-age population, and remains at rates above 20% for approximately half a century under the non-adaptive policymaker. For reference, the US unemployment rate during the Great Depression peaked at either 24.9% or 22.9%, depending on whether people with 'work relief' jobs are counted as being unemployed [47]. Given the composition of the labour force across different sectors in the DSK model as it currently stands - the agricultural sector, comprising 27% of Global employment in 2019 [48], being absent from the model - it would be unwise to consider DSK unemployment rates as equivalent to historical, real-world datasets. Moreover, contributors to the DSK model have noted that the DSK model may be better suited to the qualitative comparison of different scenarios than to quantitative predictions [27]. With these comments in mind, a note can be made of the nonetheless remarkably high unemployment rates seen, before turning to a more detailed analysis of the behaviours seen.

The unemployment rate increases with the introduction of the carbon tax, for all policymakers, under all ECSs. Whenever the carbon tax is a fixed value, in real terms, unemployment peaks when the tax is first introduced. This results from the fact that this is the time at which the economy is most fossil-fuel dependent, as evidenced by the electrification fraction and energy mix; consequently, the costs to employers - the energy sector and the capital and consumer goods manufacturers - are greatest, leading to a greatly depressed demand for labour. Once the green transition is underway, the unemployment rate under non-adaptive policymakers gradually decreases; the slower the transition, however, the slower the decrease in unemployment. It is for this reason that when the carbon tax is fixed at a higher value, as opposed to remaining constant at a lower value, the unemployment rate decreases significantly more rapidly. Additionally, there is a visible increase in the rate of decrease in unemployment in 2060 for the fixed, risk-averse policymaker, and after 2100 for the risk-neutral policymakers. This is likely because these dates correspond with the time that these policymakers achieve 100% green electricity generation - consequently, consumption good firms and electrified capital good firms are no longer affected by the carbon tax from this point on, accelerating the recovery in employment. Finally, an additional factor contributing to the faster decrease in unemployment after the initial peak when higher, constant, carbon taxes are imposed is that in the period during which the carbon tax affects firms the most, early in the transition, firms in both sectors will place a greater emphasis in investing in increasing the energy efficiency of their operations than in their labour productivity. Consequently, when production begins to recover from the initial economic shock, there will be a greater demand for labour for the same amount of production than there would otherwise be than there is in the lower-tax scenarios.

Comparing the adaptive and non-adaptive, risk-averse policymakers, the differences in unemployment are considerable. The unemployment rate under the adaptive policymaker initially follows that of their non-adaptive counterpart, before more rapidly falling to rejoin the risk-neutral rates in the late 2070s. This evolution, relative to the other policymakers, very closely follows the evolution of the carbon tax, suggesting that in the DSK, employment can quickly recover from the economic shock incurred by high carbon taxes. Nonetheless, there is a delay between the carbon taxes of the adaptive, risk-averse policymaker and the risk-neutral policymakers becoming virtually equal in 2060 and the unemployment doing the same approximately 15 years later, even if the discrepency during this period is barely a 2% difference. More generally, it can be noted that while an adaptive, risk-averse policymaker under an ECS which is in line with our current expectations will allow warming to increase significantly, relative to their non-adaptive counterpart, the gains in employment are also highly significant. For approximately 40 years of model run time, the difference in the expected unemployment rates of these two policymakers is in the 10% range.

4.2.2 Higher ECS case

Figure (4.7) shows the evolution of the (a) the carbon tax and the resulting (b) green electricity share, (c) electrification, (d) GDP, (e) CO_2 emissions, (f) temperature change and (g) unemployment under all four policymakers implementing P1 under a much higher ECS of 5.81°K/doubling CO_2.

Observing first panel (a), in this case the adaptive, risk-neutral policymaker's carbon tax increases between 2020 and 2060 to match the tax imposed by the risk-averse policymakers; this is a result of the fact that 5.81°K/doubling CO_2 is very close to the 99^{th} percentile of the Sherwood et al. PDF [22]. Also of interest in panel (a) is the behaviour of the adaptive, risk-averse policymaker's carbon tax, which resembles the evolution of the 99^{th} percentile value of the ECS PDF in figs. (3.2) and (3.3). When the model's climate and bayesian learning modules begin to run in 2020, the rate of warming is much higher than expected; however there is not enough warming data at this point to shift the ECS estimate with much accuracy, leading to an overshoot. Conversely, the research PDFs linearly approach the scenario's true value from the initial PDF, narrowing as they go - in this case, approaching the true ECS from below. This is the reason why the carbon tax subsequently undershoots that imposed by the non-adaptive policymaker in the 2030s: this follows the incorporation of the research module's first simulated ECS PDF into the policymaker's understanding of the ECS, shifting the value downwards.

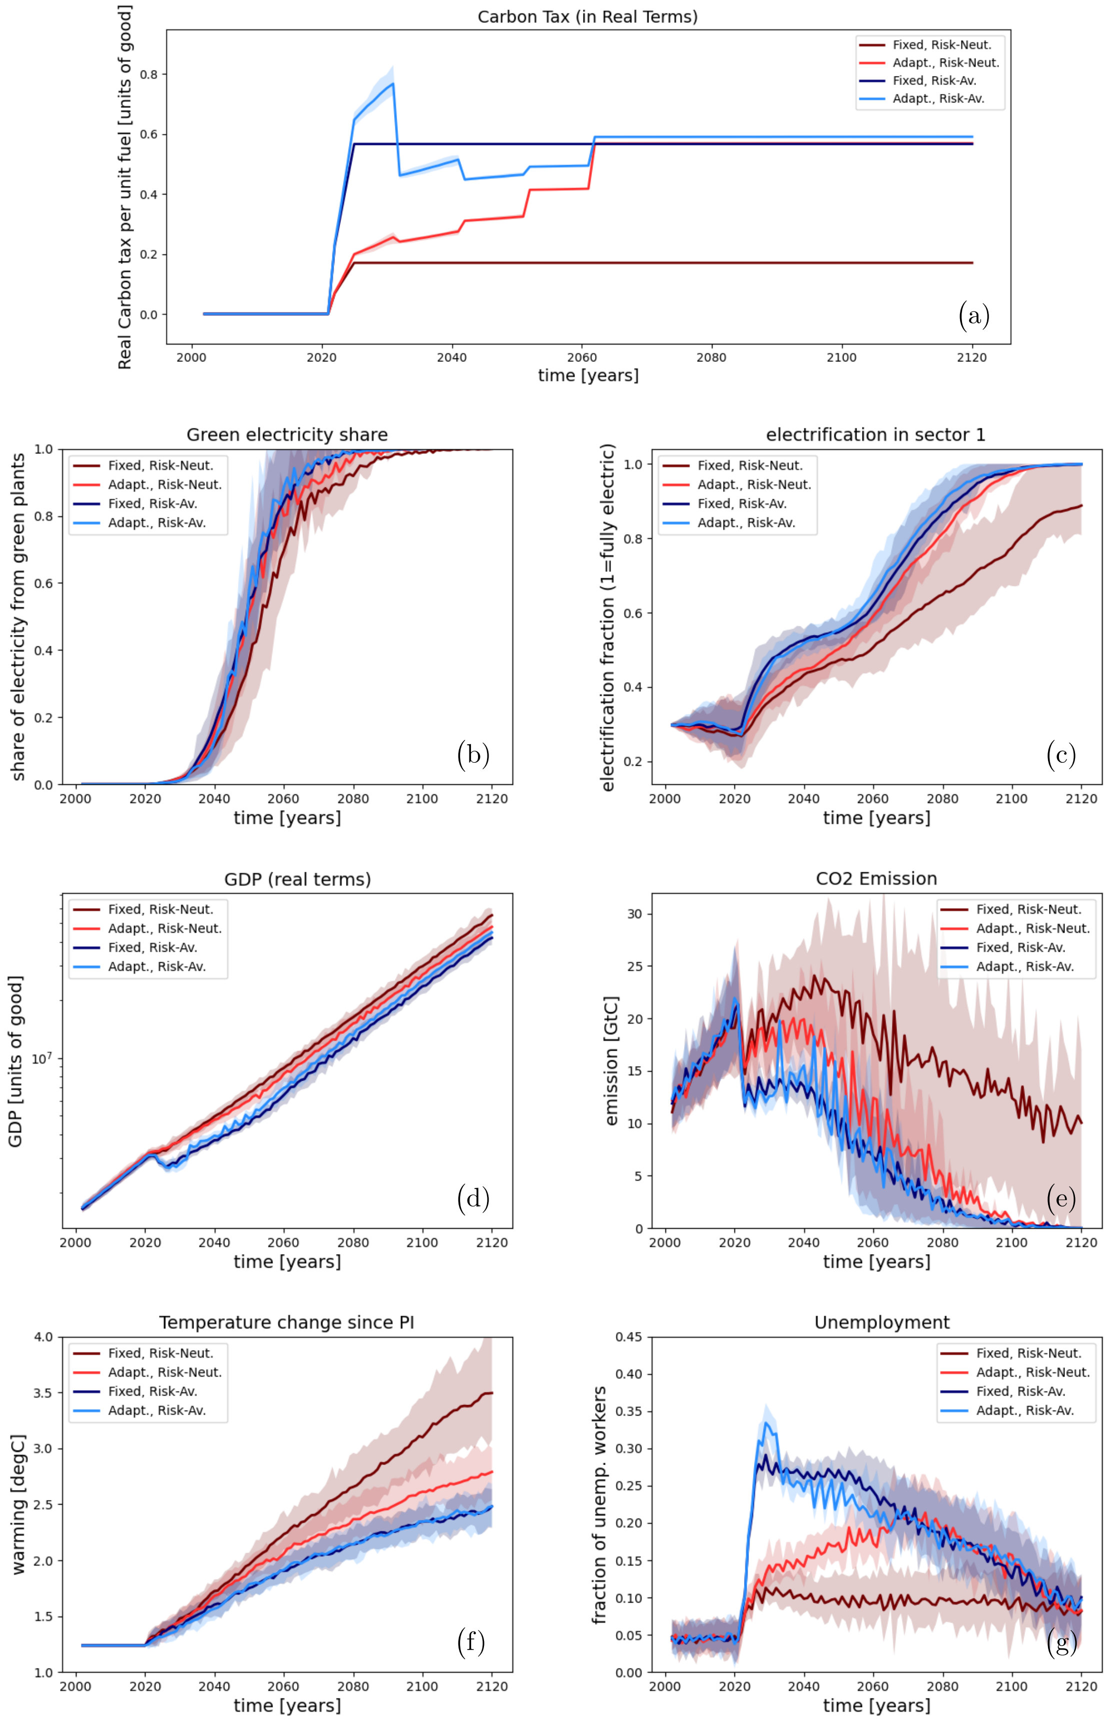

Figure 4.7: Evolution of the climate and economy from 2000 to 2120 when P1 (carbon tax only) is pursued, if ECS=5.81°K/doubling CO_2. The model has been run for 50 realisations; the shaded region in each plot corresponds to the values bounded by the 10^{th} and 90^{th} percentiles, while the dark line shows the mean time series. Shown are (a) the global mean surface temperature change with respect to pre-industrial levels, modelled from 2020; (b) the unemployment rate; (c) the share of electricity produced by green plants; (d) the fraction of the capital goods sector which has electrified; (e) the GDP, adjusted for inflation, and (f) annual CO_2 emissions.

Observing panel (b), the same behaviour noted when discussing the green electricity share in the context of the lower ECS can be seen. Higher carbon taxes tend to lead to a faster transition to 100% green electricity. However, the policymaker who imposes the highest initial carbon tax, due to the overshooting in the learning process, has the third slowest increase in green electricity share until the year 2040. Nonetheless, when the electrification is considered, the higher carbon taxes always seem to be associated with greater rates of electrification. Notably, under this ECS, all but the risk-neutral, non-adaptive policymaker 100% electrification prior to 2100. Panel (e) confirms that all but this policymaker reach zero emissions prior to 2100 in most realisations as well. Notably, observing panel (d), the later increase of the carbon tax once the energy transition is well underway in the mid century under the adaptive, risk-neutral policymaker does not cause any recession, although the GDP increases less than it would if the tax were kept lower.

Undrr this ECS, the adaptive and non-adaptive risk-averse policymakers hold the rise in temperature between 2°K and 2.5°K in approximately 80% of the realisations, although it is unclear whether the climate has stabilised at the end of the model run. This apparent failure to stabilise the climate by 2120 is evidently not the result of emissions, since both of these policymakers have emitted no CO_2 in the 20 years preceding this, as already noted. Recalling section 2.2.2, the reason for this slower stabilisation is the fact that, for higher ECSs, the TCR has been adjusted accordingly so it is much lower. The temperature therefore takes much longer to respond to the increase in atmospheric CO_2 concentrations. The risk-neutral, non-adaptive policymaker fares most poorly in climate mitigation, as might be expected. Under this policymaking strategy in the higher ECS case, temperature by 2120 reaches between 3°K and 4°K above pre-industrial values. By contrast, the adaptive, risk-neutral policymaker under the same ECS keeps temperature rises under 3°K in over 90% of cases, though the Paris Climate Accord goal of keeping warming under 2°K is far from site.

Considering panel (g), the evolution of the unemployment as a result of the carbon tax is similar to that in the lower ECS case. The unemployment is significantly different between the adaptive and non-adaptive, risk-averse policymakers in the first decade, when the carbon taxes imposed by each differ; however, the unemployment quickly drops along with the carbon tax under the adaptive, risk-neutral policymaker. Finally, the red line in panel (b) shows the case of an adaptive, risk-neutral policymaker who must adapt to a much higher-than-expected ECS. As might be expected from the observations made thus far, the unemployment rate rises as the carbon tax increases to match that of the risk-averse policymakers. Notably, however, the unemployment under this policymaker is consistently higher that that seen under the two risk-averse policymakers during the 2070s, despite the fact that the carbon tax is virtually identical under these three policymakers from 2060 onwards. This may be attributed to the fact that during this period, the energy transition and electrification rates under the adaptive, risk-neutral policymaker lag behind that of the other two, making the same carbon tax more costly to firms than under risk-averse policymakers.

4.2.3 Performance over all ECSs tested

Having analysed in detail the time evolution of the modelled global temperature and unemployment rates under the different policymakers under two ECSs, the outcomes of different policymaking strategies across *every ECS tested* are now considered. Figure (4.8(a)) shows the peak temperature change, relative to the pre-industrial value, prior to 2100. The peak before 2100, rather than over the course of the whole model run, is chosen as the focus partly because 2100 corresponds to the end of the IPCC's 'long term' timescale [49]. An additional reason is that, as noted in chapter 2, the DSK model does not feature negative emissions, which essentially assumes that no new negative emissions technologies will not be developed or scaled up over the course of the model run, and that land use will not change to increase the biosphere's ability to sequester carbon. As model time increases, this assumption becomes increasingly hard to justify; thus, 2100 is chosen as a cut-off date for judging the environmental outcomes of policies.

Observing fig. (4.8(a)), we see non-adaptive policymakers experiences greater warming for higher ECSs. Adaptive policymakers are more successful than their non-adaptive counterparts under high ECSs, but adopt less stringent mitigation strategies in lower-ECS cases, thereby seeing considerably higher warming in these situations. Recalling our discussion of P1 under an ECS of 2.99°K/doubling CO_2, only the non-adaptive, risk-averse policymaker is able to consistently reach zero emissions in this situation. This, combined with the fact that the peak temperature takes longer to be reached under higher ECSs, due to the adapted TCR, explains why the highest warming under the adaptive, risk-neutral policymaker is seen for the lowest ECS. It is at this ECS that the electrification rate is lowest and production highest, resulting in the highest emissions of any of the simulations. Overall, then, whereas there is a clear gradient of peak temperature as a function of ECS for non-adaptive policymakers under P1, this is not the case with their adaptive counterparts.

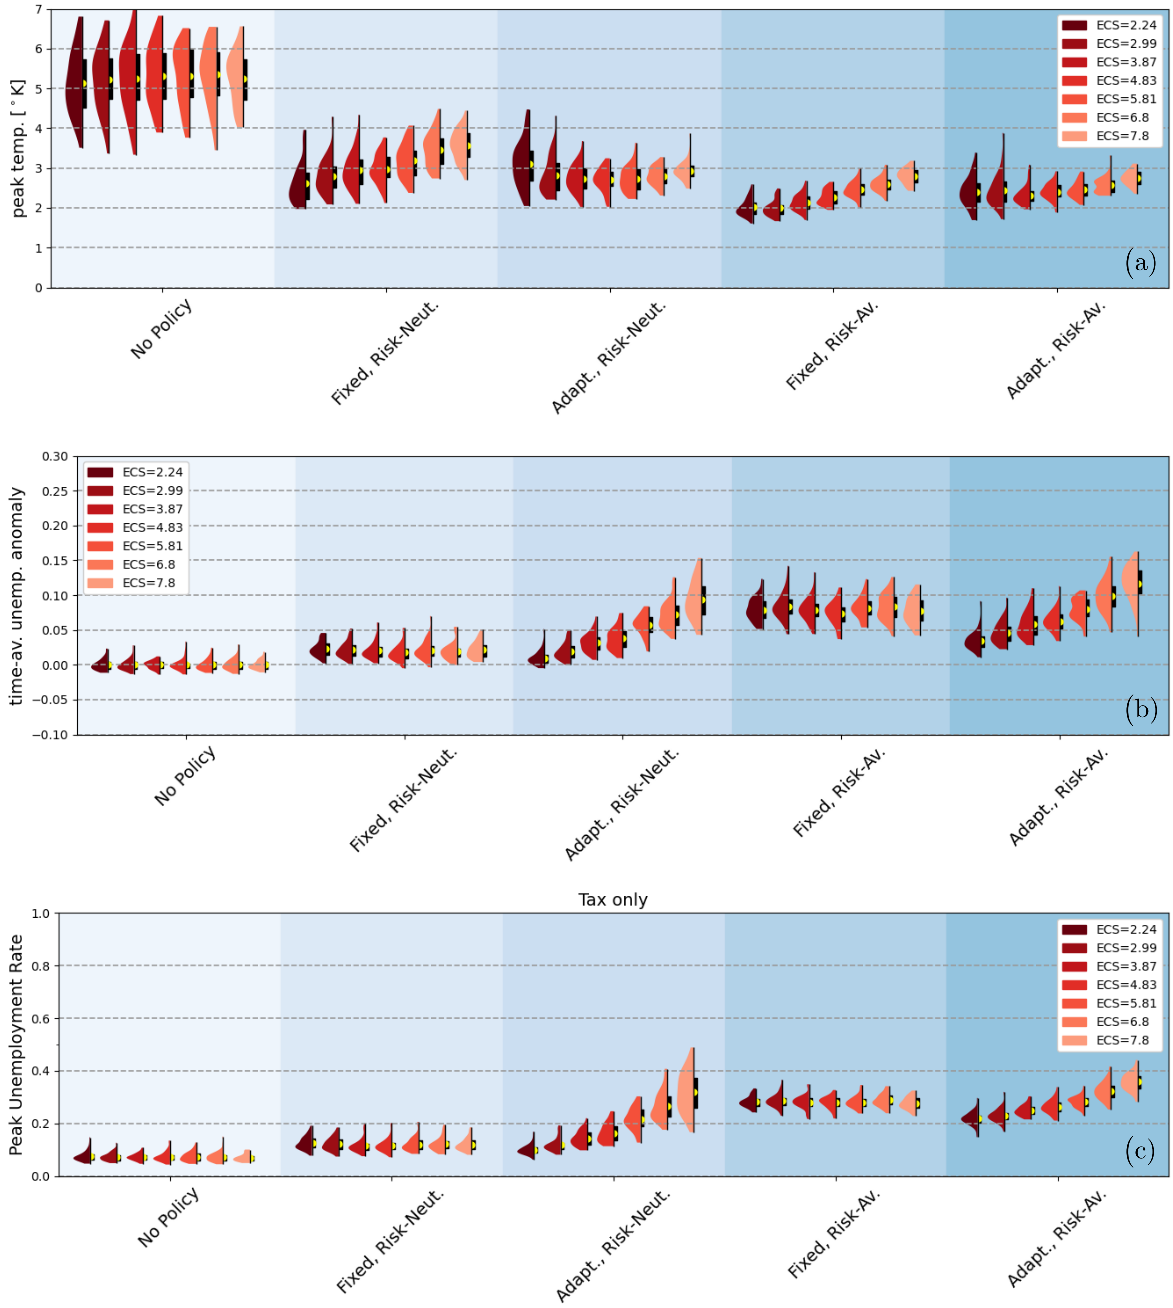

Figure 4.8: (a) Peak warming prior to 2100, relative to pre-industrial value, for all policymakers, over all ECS values tested. The policy applied in this figure is P1 (carbon tax only). As 50 realisations of the model were run for each ECS value, for each policymaker, the distribution of peak temperatures is shown for each case, in varying shades of red. For each policymaker, labelled on the x-axis and demarcated in the figure by varying shades of blue, the mean value (yellow dot), first and third quartiles (thick vertical black line) and distribution have been shown. For each policymaker, the values are plotted from left to right, sorted by increasing ECS, with the darkest shaded distribution, on the left, corresponding to the lowest tested ECS value. The ECS corresponding to each shade of red is shown in the legend, in units of °K/doubling CO_2. (b) Time-averaged unemployment difference from 2020 to 2100, where the baseline unemployment is the mean unemployment in the case where no policy is applied. The policy applied in is P1 (pure carbon tax). The values are stated in units of the fraction of the total working-age population. The plot is organised by policymaker and ECS in the same way as in panel (a). (c) Peak unemployment under P1. This is not plotted relative to the no-policy baseline due to the variability of the unemployment in the baseline case.

Across all ECSs, the spread of the peak temperatures over different model realisations is narrower the more stringent the climate policy applied. This is most visible from the difference between no policy, with different realisations having peak temperatures spread over more than 3°K, and the fixed, risk neutral policymaker, with realisations diverging by 2°K at most. This trend continues in the jump to the risk-averse policymakers, who tend to see less than a degree's divergence between different realisations. This is likely a result of the fact that variations in emissions time series between different realisations will have a greater effect on the temperature if zero emissions is never reached.

Overall, the fixed, risk-averse policymaker has the greatest chance of keeping warming to 2°K or less, given that the other policymakers rarely achieve this. While even the fixed-risk averse policymaker's performance with respect to this goal may seem poor, it should be recalled from the ECS PDF shown in fig. (3.4) that the first 3 ECSs represent approximately 90% of the probability of the PDF. Nonetheless, considering the distribution of temperature in the different realisations, and the higher ECSs, there is clearly an over 50% chance that this policymaker sees a peak temperature of over 2°K.

Panel (b) of (4.8) presents the time-averaged difference in unemployment between each policy and the expected no-policy result. By definition, non-adaptive policymakers do not change the carbon tax depending on the ECS, so it is as expected that the unemployment rate has no dependence on ECS for these policymakers. Nonetheless, observing the distribution of average unemployments for a given non-adaptive policymaker in fig. (4.8(b)) does make it possible to see that even running 50 realisations per scenario is not enough to capture all visible variation in the DSK's economic model. As noted in section, among non-adaptive policymakers, the risk averse policymaker causes significantly more unemployment than the risk-neutral policymaker, increasing time-mean unemployment by an average of 8-9%, relative to the baseline, rather than 2-4%.

For the adaptive policymakers, unemployment increases greatly with ECS, with expected values ranging from 1% for the lowest ECS to nearly 10% for the highest in the case of the risk-neutral policymaker. The spread of time-mean unemployment also increases with the ECS. This could be a result of greater variations in the evolution of the policymaker's knowledge of climate sensitivity, as it needs to increase over a much larger range than that initially expected. The risk-averse policymaker does not see as great an increase in time-mean unemployment, relative to their risk-neutral counterpart, than the non-adaptive, risk-averse policymaker does. This is likely because the adaptive, risk-averse policymaker tends to successfully reach zero emissions under higher ECSs. Consequently, the unemployment rate will recover towards the end of the model run, as the carbon tax no longer contributes to firms' costs.

Observing panel (c), the behaviour of the peak unemployment follows much of the same behaviour as the time-averaged values. However, perhaps counter-intuitively, the peak unemployment rate is at times higher for the adaptive, risk-neutral policymaker than the adaptive, risk-averse policymaker in the highest ECS case. In fact, this is a result of the Bayesian learning process in the 8 years before the first research-function-imposed narrowing of the PDF. The expected ECS of the risk-neutral policymaker under an ECS of 7.8°K/doubling CO_2, having started near 3°K/doubling CO_2 increases extremely rapidly as a results of the much higher than expected warming. This is so much the case that it is at times higher that the 99$^{\text{th}}$ percentile ECS of the risk-averse policymaker, who started with an already high ECS, in the first few years of learning. More generally, it can be noted that while the adaptive, risk-neutral policymaker sees warming of over 3°K under the highest ECS, they cause a peak in unemployment of over 40% in some realisations. This is a testament to the inefficacy of P1 as a climate policy.

4.3 P2: Carbon Tax, 50% Directed towards Green Plant Building

4.3.1 Temperature change

The climate mitigation achieved with P2 is much greater than that with P1: among the risk-neutral policymakers, the highest peak warming seen under P2 is one degree lower than under P2. In all but a few realisations under every non-baseline policymaker, warming is kept to under 3°K under every ECS. The spread of the peak temperature seen under each policymaker, under each ECS, is narrower than it was under P1. This likely indicates that zero emissions are consistently achieved by all policymakers. Risk averse policymakers have an over 90% chance of keeping peak warming below 2°K. Unlike under P1, the gradient of peak temperature as a function of ECS seen under non-adaptive policymakers is also seen under adaptive policymakers, although it is not as pronounced. The adaptive policymakers generally see under lower warming under high ECSs and higher warming under low ECSs than their non-adaptive counterparts. Nonetheless, peak warming under the adaptive, risk-averse policymaker is barely higher than under their fixed counterpart.

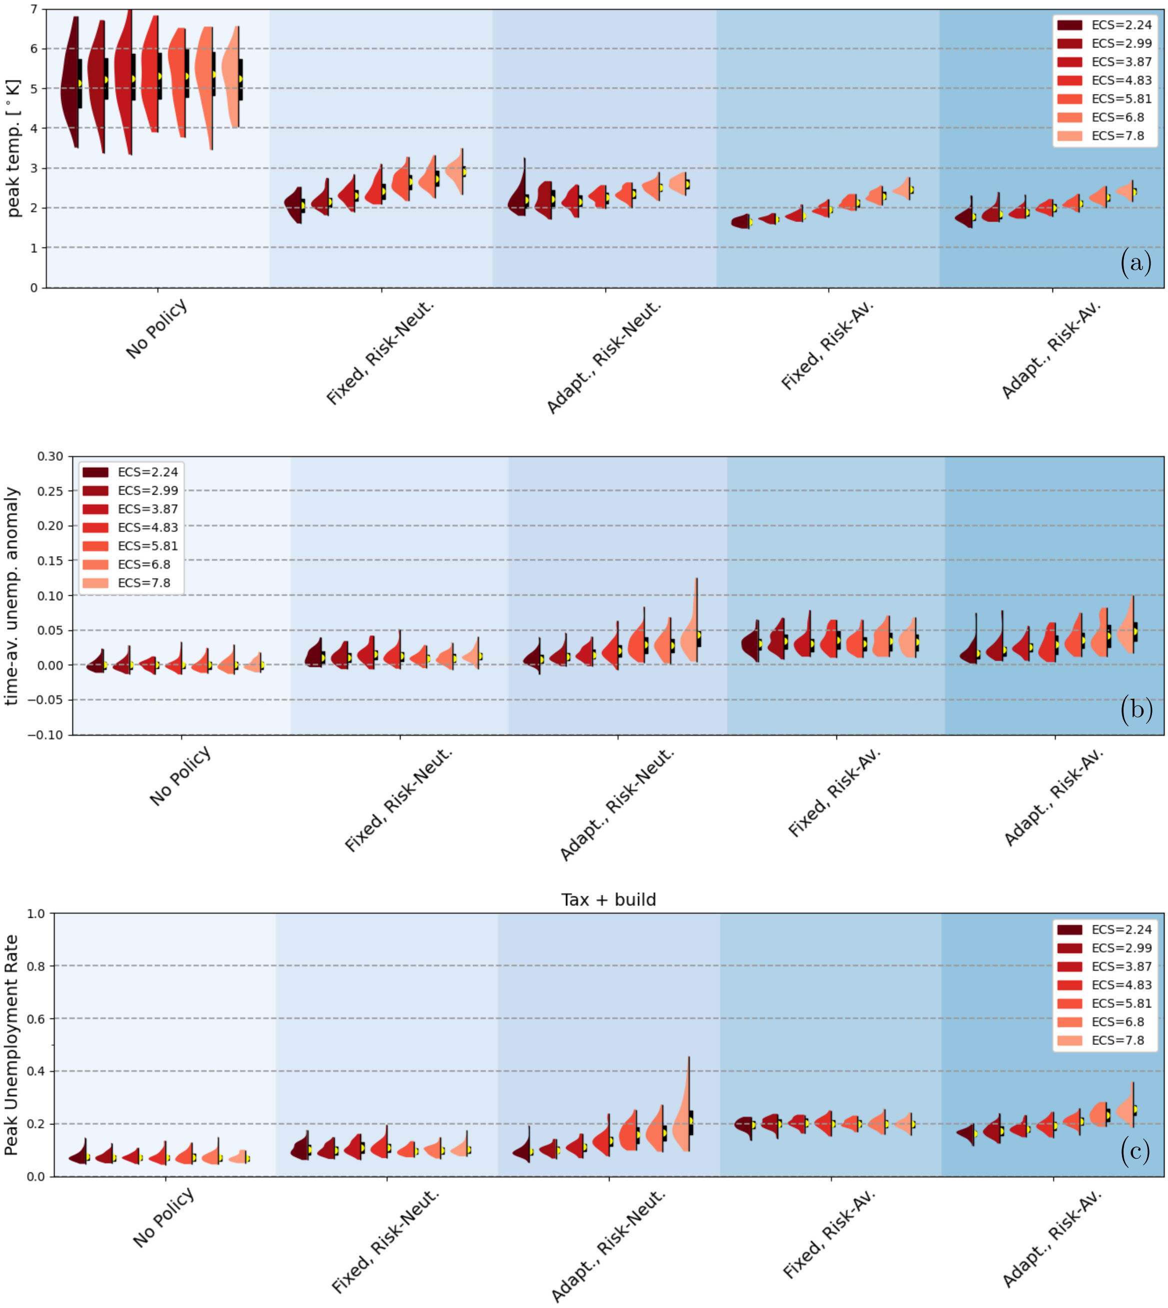

Figure 4.9: (a) Peak warming prior to 2100, relative to pre-industrial value, under P2 (carbon tax with state green plant building). The plot is organised by policymaker and ECS in the same way as in fig. (4.8). (b) Time-averaged unemployment difference from 2020 to 2100, relative to the mean unemployment in no policy case, under P2 (carbon tax with state green plant building). (c) Peak unemployment under P2.

4.3.2 Unemployment

Observing panels (b) and (c) of fig. (4.9), it is immediately clear that P2 causes less unemployment than P1; indeed, some realisations with the risk-neutral policymakers see less time-averaged unemployment than the expected baseline value. The jump in unemployment from risk-neutral to risk-averse policymakers is much smaller than under P1. The increase in time-mean unemployment with ECS under the adaptive policymakers is also much less steep than under P1. This suggests that, under P2, the time-mean unemployment rises less decisively with higher carbon taxes than under P1. This is difference is also seen in the peak unemployment, although to a lesser degree, which suggests that the gains in employment are largely due to a faster recovery from the initial spike in unemployment than under P1.

The improved performance, as gauged by unemployment, of P2 compared with P1 - as well as the fact that time-mean unemployment increases less with higher carbon taxes - can be understood as the result of two main causes. First is the fact that, under this policy, the government directly creates jobs in green plant building and maintenance: crucially, the number of jobs created is directly tied to the revenue from and hence value of the carbon tax. This is expected to slightly reduce the peak . Notably, in cases where the carbon tax is extremely high, more green plant jobs might not necessarily be created than would be by a lower carbon tax: beyond a certain level the higher carbon tax mostly causes a deeper recession, with the depressed economic activity resulting in lower carbon tax revenue and hence fewer green plant jobs than might otherwise be created with a slightly lower carbon tax.

The second, less direct, cause of the gains in employment is the faster green transition seen under this policy. A high carbon tax - although not as high as in the extreme case discussed - results in a large green plant building fund, and hence a very fast increase in the green share of electricity during the years when little green plant building takes place under P1. The experience gained in green plant building leads to a lowering of the costs of green plant building through innovation. Consequently, the rate of green plant building can be maintained at a very high level, meaning that a full transition in electricity generation can be achieved much more quickly than under P1. Consumption good firms, therefore, sooner benefit from decreased costs from the carbon tax. The sooner date of the full transition in electricity generation means that electrification becomes much more attractive to capital good firms much more quickly than would be the case under P1. The result is that the industrial sector sooner finds carbon taxes less costly than under P1, thereby sooner the sector's capacity to (re)hire.

Considering the two lowest ECSs - accounting for 71.9% of the ECS PDF - it can be seen that, despite the very minimal nature of the increased warming under the adaptive, risk-averse policymaker relative to their non-adaptive counterpart, the gains in employment are tangible. Time-mean unemployment in these cases decrease relative to the fixed policymaker by 1-2%, and the expected peak unemployment is also several percentage points lower. This, combined with the slightly lower warming in the highest ECS cases, points to a clear benefit to adaptive, risk-averse policymaking under P2.

4.4 Comparing All Policymakers

Figure (4.10) shows the effects on global warming and time-averaged unemployment of each policymaker following each policy, aggregated over every ECS tested. The figure effectively presents a summary of the main points discussed in the preceding section, but with a more formal method of considering the results over different ECSs. Temperature change and unemployment values were aggregated over different ECSs using a simple probability-weighted sum, thus yielding the expectation values for the two, following the method described in chapter 3.

Risk-neutral policymaking leads to higher expected warming than risk-averse policymaking. This difference is significant for both policies, but more so for P1. Among risk-neutral policymakers, adaptive policymaking leads to slightly higher warming, on average, than non-adaptive policymaking, a result of the much higher probability of low ECSs. Expected unemployment is also very slightly slightly higher under the adaptive policymaker when a risk-neutral approach is followed. Considering risk-averse policymakers, fixed policymaking leads to much lower expected peak temperatures than adaptive policymaking; this comes at the cost of significantly higher unemployment. This contrast is particularly pronounced under P1. Under P2, although the expected warming increases in the shift from non-adaptive to adaptive policymaking, it is still has an approximately 90% chance of being under 2°K.

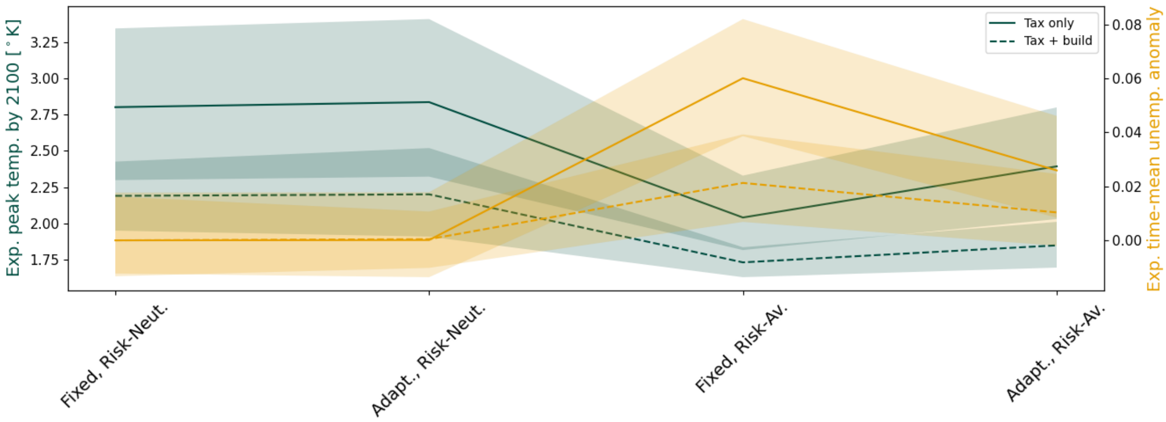

Figure 4.10: Peak temperature change, relative to pre-industrial value, and time-integrated unemployment difference, relative to null policy case, for each policymaker and policy, aggregated over all ECS values. The peak temperature change is plotted in green, and measured on the left-hand y-axis, while the unemployment is plotted in orange, with its scale on the right-hand y-axis. The values for each case have been aggregated over all ECS values tested through a weighted average, where the weighting is the probability of each ECS, according to the PDF in [22]: the temperature and unemployment values may thus be treated as estimated expectation values, according to recent estimates of the ECS. Shaded regions correspond to the values bounded by the 10^{th} and 90^{th} percentiles of the realisations simulated, with the lines corresponding to the means. As shown in the legend, solid lines correspond to policymakers pursuing P1 (carbon tax only), with dashed lines corresponding to the case where P2 (carbon tax with state green plant building) is pursued. The null policy case has been omitted to show the differences between the outcomes of the different policies and policymakers in greater detail.

Chapter 5

Discussion

It seems worthwhile to preface the conclusions drawn in the next chapter with a discussion of the strengths and weaknesses of this work's approach and methods, the model used to undertake the research, and, more generally, the difficulties in comparing outcomes of different policies.

It is worth noting that the results discussed in this report are very much dependent on the year in which this research was undertaken. While there may be some benefits to adaptive policymaking - this will be commented on in the final chapter - there would likely be more potential benefits to this approach if it had been considered in the past, at a time when the window for effective climate action was more open. Indeed, the UNEP's emissions gap report 2022, already cited several times in this work, referred to itself as a 'testimony to inadequate action on the global climate crisis' [6]. If current global trends in enacted climate policy continue, it is likely only a matter of time until the required rate of mitigation is so extreme that any questions of ECS-adaptive policymaking become largely irrelevant. Of course, it is impossible to determine *a priori* when, if ever, our research questions will become truly irrelevant - this might be, in itself, a question worthy of further study.

Focusing more concretely on the methods used in this work, a key criticism that can be made is the lack of non-carbon tax policies investigated in answering the research questions. As has been noted in the results, the differences between different policymakers are partly dependent on the policy being enacted: it would therefore be greatly beneficial to consider policies that are not reliant on a carbon tax, as the dynamics could be markedly different. This omission is largely due to time constraints, combined with the more complicated nature of making such policies adaptive. Nonetheless, one example of a policy that could be investigated would be an adaptive green plant building program, funded by time-constant, higher corporation tax. Corporation tax is suggested rather than income tax mainly because workers in the DSK model are modelled as a homogeneous mass, so the effects of differing personal income taxes would be much harder to model. Another direction further investigation might take would be the incorporation of regulation into the policies investigated. While it is hard to see how regulation could be made adaptive, it may be that the interaction between regulations and other, adaptive, policies could lead to dynamics that are not seen in this report.

Considering the climate modelling used in this work, there are also criticisms that can be made. Most pressing is the treatment of the ECS and TCR: while the approach followed was to adapt the TCR to the ECS so as to make the evolution of the GMST consistent with 2020 values, it may be that this misses a key dynamic linking the two measures of climate sensitivity, calling into question how realistic the climate modelling is. To some extent, this criticism can be mitigated by noting that the potential lack of realism is largely a reflection of the limitations of working with relatively simple reduced-complexity climate models, and that the purpose of this research is to analyse the outcomes under different policymaking approaches comparatively, rather than to produce precise quantitative predictions.

A note has to be made at this point that the main priority before undertaking further research using the models adapted in this work would be to correct the error in the implementation of the bayesian learning routine.

5.1 The Model and Wider Environmental Context

The DSK model has been noted in the context of an increasingly large base of new climate-economic models which move away from neoclassical economic assumptions, particularly that of economic equilibrium

[50]. There are benefits in using the DSK model, rather than simpler and/or more conventional IAMs. A key advantage over the DICE model, for instance, is the fact that investing in mitigation involves investment in an explicitly modelled stock of green plants, in the context of an electricity firm with its own dynamics. This allows for more detailed conclusions than would a similar exercise conducted using DICE. The fact that electrification, as well as a transformation of the energy mix, is necessary for full decarbonisation is an additional benefit. Finally, the agent-based nature of the model's industrial sector means that a rise in unemployment following the introduction of a large carbon tax emerges from the costs incurred on individual employers, rather than from a set of theoretical macroeconomic relations, to some extent limiting the number of assumptions and simplifications the model has to make.

Naturally the DSK model is not without its own shortcomings; some of the more pressing ones for this research can briefly be considered. In some ways the model is overly pessimistic, from a climate mitigation perspective: there is no limit to the amount of fuel that can be burnt; this, combined with the perhaps overly robust annual GDP growth rate of approximately 3%, may result in too much warming in the no-policy baseline case. On the other hand, the model does not currently feature any negative environmental or social consequences of ever increasing energy use. There is no limit to the number of green plants that can be built; this may be problematic from the perspective of the ecological and health impacts of the extraction of the materials used to build them. While it may be difficult to explicitly model these risks in the DSK model, perhaps additional ecological indicators could be incorporated into the model as well.

5.2 Weighing up different policies

Contrary to the title of this work, no strict optimisation method is followed to find the preferred policy. This is partly a result of the technical infeasibility of such an approach using the DSK model, due to its computationally intensive nature. The result has been a more explicit discussion of the trade-offs and synergies found between different policies and strategies. However, this has also led to a certain difficulty in determining how to weigh up different social and environmental ills, especially considering that the two are often interconnected. Indeed, considering the topic of social ills, unemployment is used in this work as the sole indicator of negative social consequences. It should be noted, though, that some climate economists see a society wide decrease in total working hours - in effect reducing or eliminating unemployment through a wide decrease in the working week - as a potentially positive outcome of thorough climate mitigation [51]. Whether or not this is the case, however, there is a more general question of how to weight social and environmental ills - even if this weighting is more conceptual than strictly mathematical. This project has attempted to leave the question as open as possible, and leave the reader to come to their own conclusions, perhaps in the spirit of reaching such decisions through discussion.

Chapter 6

Conclusion

This work aimed to answer two primary research questions:

1. How aggressive should policymakers be in their approach to mitigation while the Equilibrium Climate Sensitivity (ECS) is still largely uncertain?
2. Should policymakers adapt their strategy as understanding of the ECS evolves?

To this end, an agent-based integrated assessment model, the DSK model, has been adapted to incorporate policymakers who learn the climate sensitivity as the global mean surface temperature increases. The outcomes, in global warming and unemployment, seen under these policymakers are compared with those seen under policymakers who do not adapt their approach as new information on the ECS is made available. Similarly, risk-neutral policymakers, who are concerned only with the expectation value of the ECS are compared with risk-averse policymakers, who act according to the evolution of the 99^{th} percentile value of the ECS. The differences in outcomes under these four policymakers are discussed in the context of two different policies, reflecting the fact that the best strategy may depend on the type of policy pursued.

Our conclusions will rest on the observations made in section 4.4. In terms of coming close to meeting the Paris Agreement goals, it is clear that, under both policies investigated, only risk-averse policymaking is effective. However, within this category the conclusions become more nuanced. When the results are aggregated across different ECSs according to current understanding of their approximate probability, non-adaptive policymaking achieves the greatest climate mitigation. This also comes at the cost of higher unemployment, even under the second policy, in which the policymaker uses the revenues of their carbon tax to build green plants. To the author of this report, the outcomes of the adaptive, risk-averse policymaker implementing the second policy seem the most attractive. What is clear, however, is that whatever the relative merits of adaptive and non-adaptive policymaking, the choice of which policy to pursue has considerably more impact, in terms of both unemployment and climate change. In conclusion, then, risk-aversion regarding the ECS is likely to lead to better outcomes and there are benefits associated with adaptive policymaking, but most important is the set of policies implemented, and at this time it is thus crucial that more policies than ever are considered, modelled, and discussed in society at large.

Bibliography

[1] S. Sarkar, "Pakistan floods pose serious health challenges," *BMJ*, vol. 378, 2022. [Online]. Available: https://www.bmj.com/content/378/bmj.o2141

[2] F. Fazel-Rastgar and V. Sivakumar, "Weather pattern associated with climate change during canadian arctic wildfires: A case study in july 2019," *Remote Sensing Applications: Society and Environment*, vol. 25, p. 100698, 2022.

[3] D. R. Fisher and S. Nasrin, "Climate activism and its effects," *Wiley Interdisciplinary Reviews: Climate Change*, vol. 12, no. 1, p. e683, 2021.

[4] C. A. Ogunbode, R. Doran, and G. Böhm, "Exposure to the IPCC special report on 1.5°C global warming is linked to perceived threat and increased concern about climate change," *Climatic Change*, vol. 158, no. 3, pp. 361–375, 2020.

[5] M. S. Schäfer and J. Painter, "Climate journalism in a changing media ecosystem: Assessing the production of climate change-related news around the world," *Wiley Interdisciplinary Reviews: Climate Change*, vol. 12, no. 1, p. e675, 2021.

[6] U. N. E. Programme, *The Closing Window: Climate crisis calls for rapid transformation of societies*, 2022. [Online]. Available: https://www.unep.org/emissions-gap-report-2022

[7] N. Stern, *The economics of climate change: the Stern review.* cambridge University press, 2007.

[8] F. P. Incropera, *Climate change: a wicked problem: complexity and uncertainty at the intersection of science, economics, politics, and human behavior.* Cambridge University Press, 2016.

[9] A. Pettifor, *The case for the green new deal.* Verso Books, 2020.

[10] J. Hickel and G. Kallis, "Is green growth possible?" *New Political Economy*, vol. 25, no. 4, pp. 469–486, 2020. [Online]. Available: https://doi.org/10.1080/13563467.2019.1598964

[11] J. C. J. M. van den Bergh and G. Kallis, "Growth, a-growth or degrowth to stay within planetary boundaries?" *Journal of Economic Issues*, vol. 46, pp. 909–920, 12 2012.

[12] R. Mastini, G. Kallis, and J. Hickel, "A green new deal without growth?" *Ecological Economics*, vol. 179, p. 106832, 1 2021.

[13] Z. Wang, J. Wu, C. Liu, G. Gu *et al.*, *Integrated assessment models of climate change economics.* Springer, 2017.

[14] F. Ackerman, S. J. DeCanio, R. B. Howarth, and K. Sheeran, "Limitations of integrated assessment models of climate change," *Climatic change*, vol. 95, no. 3, pp. 297–315, 2009.

[15] W. D. Nordhaus, "An optimal transition path for controlling greenhouse gases," *Science*, vol. 258, no. 5086, pp. 1315–1319, 1992. [Online]. Available: http://www.jstor.org/stable/2880417

[16] F. Lamperti, G. Dosi, M. Napoletano, A. Roventini, and A. Sapio, "Faraway, so close: coupled climate and economic dynamics in an agent-based integrated assessment model," *Ecological Economics*, vol. 150, pp. 315–339, 2018.

[17] R. A. Williams, "Lessons learned on development and application of agent-based models of complex dynamical systems," *Simulation Modelling Practice and Theory*, vol. 83, pp. 201–212, 2018.

[18] K. Safarzyńska and J. C. J. M. van den Bergh, "ABM-IAM: optimal climate policy under bounded rationality and multiple inequalities," *Environmental Research Letters*, vol. 17, no. 9, p. 094022, sep 2022. [Online]. Available: https://dx.doi.org/10.1088/1748-9326/ac8b25

[19] R. S. Pindyck, "The use and misuse of models for climate policy," *Review of Environmental Economics and Policy*, 2020.

[20] IPCC, "Annex VII: Glossary [van Diemen, R., J.B.R. Matthews, V. Möller, J.S. Fuglestvedt, V. Masson-Delmotte, C. Méndez, A. Reisinger, S. Semenov (eds)]." in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge, UK and New York, NY, USA: Cambridge University Press, 2021.

[21] IPCC, "Summary for policymakers," in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, V. Masson-Delmotte, P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou, Eds. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2021. [Online]. Available: https://www.ipcc.ch/report/ar6/wg1/downloads/report/IPCC_AR6_WGI_SPM.pdf

[22] S. C. Sherwood, M. J. Webb, J. D. Annan, K. C. Armour, P. M. Forster, J. C. Hargreaves, G. Hegerl, S. A. Klein, K. D. Marvel, E. J. Rohling, M. Watanabe, T. Andrews, P. Braconnot, C. S. Bretherton, G. L. Foster, Z. Hausfather, A. S. von der Heydt, R. Knutti, T. Mauritsen, J. R. Norris, C. Proistosescu, M. Rugenstein, G. A. Schmidt, K. B. Tokarska, and M. D. Zelinka, "An assessment of earth's climate sensitivity using multiple lines of evidence," 12 2020.

[23] O. Andrews, C. Le Quéré, T. Kjellstrom, B. Lemke, and A. Haines, "Implications for workability and survivability in populations exposed to extreme heat under climate change: a modelling study," *The Lancet Planetary Health*, vol. 2, no. 12, pp. e540–e547, 2018.

[24] M. L. Weitzman, "On modeling and interpreting the economics of catastrophic climate change," *The review of economics and statistics*, vol. 91, no. 1, pp. 1–19, 2009.

[25] K. Gillingham, W. Nordhaus, D. Anthoff, G. Blanford, V. Bosetti, P. Christensen, H. McJeon, and J. Reilly, "Modeling uncertainty in integrated assessment of climate change: A multimodel comparison," *Journal of the Association of Environmental and Resource Economists*, vol. 5, no. 4, pp. 791–826, 2018.

[26] W. D. Nordhaus, "The economics of tail events with an application to climate change," *Review of Environmental Economics and Policy*, vol. 5, no. 2, pp. 240–257, 2011.

[27] C. Wieners, F. Lamperti, R. Buizza, A. Roventini, and G. Dosi, "Uniform, gradually increasing carbon tax is not an effective climate policy," 2022, manuscript in prep.

[28] J. Sterman, T. Fiddaman, T. R. Franck, A. Jones, S. McCauley, P. Rice, E. Sawin, and L. Siegel, "Climate interactive: the C-ROADS climate policy model," 2012.

[29] J. C. Van den Bergh, "Relax about gdp growth: implications for climate and crisis policies," *Journal of Cleaner Production*, vol. 18, no. 6, pp. 540–543, 2010.

[30] Gallup, "Environmental protection vs. economic growth," 2022, retrieved online from Gallup's page on environment polling, https://news.gallup.com/poll/1615/environment.aspx, accessed 29/12/2022.

[31] IPSOS UK, "The economist – economic growth," 2022, retrieved online from Gallup's page on environment polling, https://www.ipsos.com/sites/default/files/ct/news/documents/2022-09/ipsos-economic-growth-the-economist-polling-september-2022.pdf, accessed 29/12/2022.

[32] IRENA, *Renewable energy benefits: leveraging local capacity for solar PV.* International Renewable Energy Agency, Abu Dhabi, 2017.

[33] ——, *Renewable energy benefits: leveraging local capacity for onshore wind.* International Renewable Energy Agency, Abu Dhabi, 2017.

[34] ——, *Renewable energy benefits: leveraging local capacity for offshore wind.* International Renewable Energy Agency, Abu Dhabi, 2018.

[35] R. Desrochers, A. Grégoire, J. Choteau, and E. Turmel, *Job Creation and Economic Development Opportunities in the Canadian Hydropower Market.* HEC Montreal, 2011.

[36] IRENA, *Renewable energy benefits: renewable energy and jobs 2022.* International Renewable Energy Agency, Abu Dhabi, 2022.

[37] A. P. Schurer, M. E. Mann, E. Hawkins, S. F. Tett, and G. C. Hegerl, "Importance of the pre-industrial baseline for likelihood of exceeding Paris goals," *Nature Climate Change*, vol. 7, pp. 563–567, 8 2017.

[38] WMO, P. Taalas, and A. Guterres, *State of the global climate 2020.*

[39] P. Friedlingstein, M. O'Sullivan, M. W. Jones, R. M. Andrew, L. Gregor, J. Hauck, C. Le Quéré, I. T. Luijkx, A. Olsen, G. P. Peters *et al.*, "Global carbon budget 2022," *Earth System Science Data*, vol. 14, no. 11, pp. 4811–4900, 2022.

[40] R. J. Millar, Z. R. Nicholls, P. Friedlingstein, and M. R. Allen, "A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions," *Atmospheric Chemistry and Physics*, vol. 17, no. 11, pp. 7213–7228, 2017.

[41] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, "Bayesian data analysis (2nd edn)," 2004.

[42] G. Danabasoglu, J.-F. Lamarque, J. Bacmeister, D. A. Bailey, A. K. DuVivier, J. Edwards, L. K. Emmons, J. Fasullo, R. Garcia, A. Gettelman, C. Hannay, M. M. Holland, W. G. Large, P. H. Lauritzen, D. M. Lawrence, J. T. M. Lenaerts, K. Lindsay, W. H. Lipscomb, M. J. Mills, R. Neale, K. W. Oleson, B. Otto-Bliesner, A. S. Phillips, W. Sacks, S. Tilmes, L. van Kampenhout, M. Vertenstein, A. Bertini, J. Dennis, C. Deser, C. Fischer, B. Fox-Kemper, J. E. Kay, D. Kinnison, P. J. Kushner, V. E. Larson, M. C. Long, S. Mickelson, J. K. Moore, E. Nienhouse, L. Polvani, P. J. Rasch, and W. G. Strand, "The community earth system model version 2 (CESM2)," *Journal of Advances in Modeling Earth Systems*, vol. 12, no. 2, p. e2019MS001916, 2020, e2019MS001916 2019MS001916. [Online]. Available: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001916

[43] W. C. V. Etten, *Introduction to Random Signals and Noise.* Wiley, 2005.

[44] E. P. Peter E. Kloeden, *Numerical Solution of Stochastic Differential Equations*, 1st ed., ser. Applications of Mathematics 23. Springer Berlin, Heidelberg, 1992.

[45] X. Lan, P. Tans, and K. Thoning, "Trends in globally-averaged CO2 determined from NOAA Global Monitoring Laboratory measurements," 2022, version 2022-12 NOAA/GML, gml.noaa.gov/ccgg/trends/, accessed 31/12/2022.

[46] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*, 3rd ed. Cambridge University Press, 2006.

[47] R. A. Margo, "Employment and unemployment in the 1930s," *Journal of Economic Perspectives*, vol. 7, no. 2, pp. 41–59, 1993.

[48] International Labour Organization, ILOSTAT database, "Employment in agriculture (% of total employment) (modeled ILO estimate)," 2021, data retrieved from World Development Indicators, https://data.worldbank.org/indicator/SL.AGR.EMPL.ZS, accessed 09/12/2022.

[49] IPCC, "Summary for Policymakers," in *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, P. Shukla, J. Skea, R. Slade, A. A. Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, and J. Malley, Eds. Cambridge, UK and New York, NY, USA: Cambridge University Press, 2022.

[50] S. Hafner, A. Anger-Kraavi, I. Monasterolo, and A. Jones, "Emergence of new economics energy transition models: A review," *Ecological Economics*, vol. 177, p. 106779, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921800919307475

[51] L. C. King and J. C. van den Bergh, "Worktime reduction as a solution to climate change: Five scenarios compared for the uk," *Ecological Economics*, vol. 132, pp. 124–134, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921800916302579

Appendix A

Evolution of the Research PDF

From eq. (3.5), it can be found that in order to produce a lognormal distribution with the a desired mean μ_x and standard deviation σ_x, repsectively, the parameters should be chosen following

$$\begin{aligned}\mu &= \ln\left(\frac{\mu_x^2}{\sqrt{\mu_x^2+\sigma_x^2}}\right),\\ \sigma &= \sqrt{\ln\left(1+\frac{\sigma_x^2}{\mu_x^2}\right)}.\end{aligned} \tag{A.1}$$

Finally, the peak of the distribution is defined as

$$M = e^{\mu-\sigma^2}, \tag{A.2}$$

as can be verified by maximising eq. (3.4).

As stated in section (3.1.2), the researched PDFs linearly approach the true ECS, such that the peak of the distribution reaches the true value in the last PDF, while the standard deviation of the distribution reaches a value of one step in ECS ($dECS$) in the final PDF. Consequently, the peak evolves according to

$$M(t) = m_M t + M_0, \tag{A.3}$$

where, from eq. (A.2), $M_0 = M(0) = \exp(\mu_0 - \sigma_0^2)$, μ_0 and σ_0 being the parameters that define the initial PDF. At the time of the last research, T (40 years, in the DSK), $M = ECS_T$, the true ECS:

$$ECS_T = m_M T + M_0, \tag{A.4}$$

which yields an expression for m_M, which can be substituted into eq. (A.3) and combined with eq. (A.2) to give

$$e^{\mu-\sigma^2} = \frac{ECS_T - e^{\mu_0-\sigma_0^2}}{T} t + e^{\mu_0-\sigma_0^2},$$

and so

$$\mu(t) = \sigma(t)^2 + \ln\left(\frac{ECS_T - e^{\mu_0-\sigma_0^2}}{T} t + e^{\mu_0-\sigma_0^2}\right). \tag{A.5}$$

Similarly, $\sigma_x(T) = dECS$, and so

$$\sigma_x(t) = \frac{dECS - \sigma_{x_0}}{T} t + \sigma_{x_0}. \tag{A.6}$$

Note that the standard deviation of the initial distribution, σ_{x_0} can be found from μ_0 and σ_0 by combining eq. (A.1), and rearranging to get

$$\sigma_{x_0} = e^{\mu_0}\sqrt{e^{2\sigma_0^2} - e^{\sigma_0^2}}.$$

Now, by combining and rearranging eq. (A.1), we also find the expression

$$\mu_x = e^{\mu+\sigma^2/2},$$

which, when combined with eq. (A.5) gives

$$\mu_x = Me^{3\sigma_2/2}.$$

Substituting this into the definition for σ given by eq. (A.1) and rearranging, we have

$$\sigma^2 = \ln\left(e^{3\sigma^2} + \frac{\sigma_x^2}{M^2}\right) - 3\sigma^2.$$

Taking the exponential of both sides, it then becomes clear that σ can be found by solving

$$e^{4\sigma^2} - e^{3\sigma^2} = \left(\frac{\sigma_x}{M}\right)^2. \tag{A.7}$$

This is achieved numerically, using a linear bisection method. Once σ has been found, μ is found from eq. (A.5).

Due to the narrowness of the distribution once the final researched PDF has been incorporated, subsequent Bayesian steps will have little effect on the PDF after this point. Figure (A.1) shows the evolution of the ECS PDF when ECS=6°K/doubling CO_2.

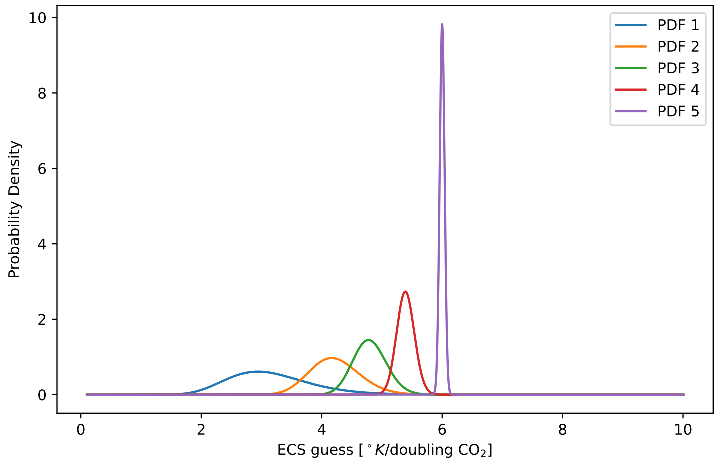

Figure A.1: Evolution of the researched ECS PDFs under a true ECS of 6°K/doubling CO_2. PDF 1 is the initial ECS PDF, which is a lognomal fit of Sherwood et al.'s estimated PDF [22].