# Uncertainty in climate response to carbon dioxide and implications for mitigation policy


**Richard James Millar**

New College


Trinity Term, 2016


Atmospheric, Oceanic and Planetary Physics

Department of Physics

University of Oxford

# Uncertainty in climate response to carbon dioxide and implications for mitigation policy

Richard James Millar, New College

Trinity Term 2016

## Abstract

Global mean surface air temperature (GMST) change, due to increases in atmospheric carbon dioxide concentrations, is a fundamental measure of human induced climate change and its associated impacts. Estimates of the magnitude of GMST response to a doubling of atmospheric carbon dioxide (equilibrium climate sensitivity or ECS) remains uncertain over a broad range between 1.5-4.5K. Expected economic damages associated with climate change are strongly sensitive to this broad uncertainty, creating challenges in constructing mitigation policies to limit peak warming.

In this thesis I explore uncertainty in the climate response through two methods. I show that potential constraints on the climate response using observations of a recent decade and planetary energy balance models are consistent with a low climate response that is not sampled by multi-model or perturbed physics ensembles of general circulation models (GCMs). I therefore subsequently set out to explore the physicality of the lower bound of climate response uncertainty in GCMs by conducting an emulator-driven perturbed physics ensemble search for low ECS models. I find a set of GCMs with ECS between 1.5K and 2K, driven primarily by negative feedbacks in tropical low cloud with GMST warming. Achieving low ECS is associated with reduced simulation fidelity relative to the standard version of the GCM, but fidelity reductions are judged to be insufficient to assuredly rule out an ECS between 1.5-2K in the real climate system.

This experiment highlights the difficulty in further reducing climate response uncertainty in the near future. I therefore propose a new framing of mitigation policy that focuses on using physical methods to index the future evolution of policy variables to emergent climate change. Such a framing could lead to adaptive mitigation policies, aimed at limiting peak warming, that are demonstrably more robust under currently irreducible physical climate response uncertainty than conventional mitigation scenarios.

# Acknowledgements

I would firstly like to acknowledge my supervisors, without whom this thesis wouldn't have been possible. I thank Professor Myles Allen for his enthusiasm, unfailing positivity, and scientific inspiration throughout this project as well as for the fascinating diversity of climate-related activities and interests that I have been exposed to under his guidance. His drive to achieve real-world impact has been a fantastic example to follow as a young researcher. I'd like to thank Dr William Ingram, for his careful guidance and patience throughout this project, which has been essential in helping me to become a better scientist over the last four years. In particular, his encyclopaedic knowledge of climate modelling and his meticulous efforts to help me slowly improve my scientific writing have been invaluable. I would also like to thank Dr Jason Lowe, for important guidance at crucial points in this project as well as helping to financially support me, along with a Natural Environment Research Council studentship, through a Met Office CASE award. I thank both of my examiners, Dr Jan Fuglestvedt and Professor Ray Pierrehumbert for a stimulating and very thoughtful discussion of this thesis in the unique location of Stockholm University.

I would like to thank the many other academic colleagues and friends that I have worked with other the last four years. Without the efforts and support of Andy Bowery, Jonathan Miller, Sarah Sparrow and the rest of the climate*prediction*.net team, many of the uniquely climate*prediction*.net problems encountered in my research would never have been overcome. The many occupants of Room 114, including Adam Levy, Niel Bowerman, Simon Driscoll, Zebedee Nicholls and Chris Huntingford have provided many fun distractions and fascinating conversations that have greatly enriched my time in AOPP. I also owe thanks to Professor Cameron Hepburn, Alex Pfeiffer and Alex Teytelboym for allowing me to pretend I know something about economics and for showing me a new world of interesting and simulating research questions. I would also like to thank the many thousands of volunteers who completed the climate*prediction*.net simulations conducted in this thesis.

I would like to thank the many friends that I have made in Oxford throughout my time as a DPhil student, both in AOPP and across the university, but particularly through New College MCR, which has provided many long-lasting friendships and a great community with which to escape department life. In particular, Elle and Rupert have been fantastic housemates over the last three years and this experience wouldn't have been the same if I hadn't shared it with them. I owe a big thank you to Belinda, whose kindness and support has kept me going throughout the long end to this thesis and has made the last year a lot more enjoyable than it might have been otherwise.

Finally, and most importantly, I would like to thank my Mum and Dad, to whom I owe more thanks than can be ever expressed here. Without their unfailing support for Robert and myself throughout our whole lives, getting here would not have been possible at all.

# List of Acronyms

| | |
|---|---|
| ALPI | Angular LTS/Precipitation Index |
| AR | Assessment Report |
| AWI | Anthropogenic Warming Index |
| BEAM | Bolin and Eriksson Adjusted Model |
| CCR | Constant Climate Resistivity |
| CDR | Carbon Dioxide Removal |
| CERES | Clouds and the Earth's Radiant Energy System |
| CMIP | Coupled Model Intercomparison Project |
| CRE | Cloud Radiative Effect |
| ECS | Equilibrium Climate Sensitivity |
| EMIC | Earth System Model of Intermediate Complexity |
| EOF | Empirical Orthogonal Function |
| ERF | Effective Radiative Forcing |
| ESM | Earth System Model |
| FAIR | Finite Amplitude Impulse Response |
| GCM | General Circulation Model |
| GDP | Gross Domestic Product |
| GMST | Global Mean Surface Temperature |
| GPCP | Global Precipitation Climatology Project |
| GTP | Global Temperature Potential |
| HadCM3 | Hadley Centre Coupled Model Version 3 |
| HLLO | High Latitude Land and Ocean |
| IAM | Integrated Assessment Model |
| IPCC | Intergovernmental Panel on Climate Change |
| IR | Impulse-Response |
| IRF | Impulse-Response Function |
| $iIRF_{100}$ | Integrated Impulse-Response Function at 100 years |
| ITCZ | Intercontinental Convergence Zone |
| LLO | Low Latitude Ocean |
| LML | Low Mid-latitude Land |
| LOO | Leave-One-Out |
| LOOT | Leave-Out-One-Third |
| LTS | Lower Tropospheric Stability |
| LW | Long-Wave |
| MAGICC | Model for the Assessment of Greenhouse-Gas Induced Climate Change |
| MLO | Mid-Latitude Ocean |
| MME | Multi-Model Ensemble |
| NWP | Numerical Weather Prediction |
| OLR | Outgoing Long-wave Radiation |
| PBL | Planetary Boundary Layer |
| PI-IR | Pre-industrial Impulse-Response |
| PPE | Perturbed Parameter Ensemble |
| RMSE | Root Mean Squared Error |
| SW | Short-Wave |
| TCR | Transient Climate Response |
| TCRE | Transient Climate Response to Cumulative Emissions |
| TOA | Top-of-atmosphere |
| UNFCCC | United Nations Framework Convention on Climate Change |
| WG | Working Group |

# Contents

# CHAPTER 1

# Introduction

## 1.1 Understanding and predicting the climate system

Climate science is both a science focused on understanding the interactions between fundamental physical processes underlying weather and climate, and a science about predicting the future.

Stretching back to Eratosthenes and Ptolemy, humanity has long made endeavours to understand the geophysical phenomena around us [Edwards, 2011]. Halley and Hadley developed the theory of how strong solar heating in the tropics combines with the fundamental laws of fluid mechanics to produce the by-then familiar system of tropical trade winds [Halley, 1686; Hadley, 1735]. Since those major insights, knowledge has been steadily improving about how individually well-understood processes interact in complex ways to determine the emergent phenomena of weather and climate.

In addition to the desire to improve understanding of the Earth system for its own sake, the improvement of future forecasts has always been a parallel motive. The initial development of wide-spread weather-observation and information networks was driven by nation-states with a desire to improve weather prediction for military and commercial reasons. In the decades after the end of the Second World War, newly available computing power (itself a product of a military desire for technology) was put to use by military research agencies to kick-start the revolution numerical weather prediction (NWP) technologies that gave rise to the subsequent progress in weather and climate prediction seen today [Edwards, 2010].

These twin motives for understanding the Earth system has led to the development of the science of *anthropogenic global warming*, namely that demonstrably significant global-scale changes in the Earth's climate system are occurring as a result of human activity. Progress in understanding the physics of the climate system has resulted in a high degree of consensus within the scientific community that human emissions from economic processes have been responsible for a high fraction of global-mean near-surface air temperature (GMST) rise observed since the mid 20$^{\text{th}}$ century [Bindoff, 2013]. Similarly, there has been substantial interest in producing projections of future climate conditions under different global anthropogenic emissions scenarios in order to inform policy, and policy-makers, about decisions regarding future energy infrastructure. Key to the accuracy of these projections is an understanding of the physical dynamics of the climate system's response when driven by anthropogenic forcing agents such as greenhouse gas and aerosol emissions. In order for uncertainties in projections of future climate change to be fully quantified and ultimately reduced, which may lead to improved decision making by policy makers [Jensen & Traeger, 2013], a better fundamental understanding of the physics of the climate system is required.

## 1.2 The development of numerical models of the climate system

One of the characteristic features of the meteorological and climate sciences is the inability to perform experiments on the scale of the entire Earth. Other branches of the physical sciences have conventionally progressed by reducing the phenomena of interest to smaller and more controllable physical systems, on which laboratory experiments can be performed and from which conclusions can be drawn to explain observations of uncontrolled natural conditions. Earth system science does not have the same luxury. To understand the complex phenomena of the planet it is necessary to understand how physical process, occurring over spatial scales ranging from a few microns to the entire size of the planet, interact. In addition to these spatial scale challenges, the temporal scales of processes that influence global climatic behaviour extends from minutes and hours through to millennia. As it is impossible to build a laboratory system of equal complexity to the natural Earth system, it is necessary to build a hierarchy of models of the Earth system at varying levels of complexity that aim to capture the required complexity needed to understand the behaviour observed in the real climate system.

This thesis will use two kinds of models to represent the climate system, general circulation models (GCMs) and models of reduced complexity. They are introduced in section 1.2.1 and section 1.2.2 respectively.

### 1.2.1 General circulation models

The most complex and comprehensive of the hierarchy of climate system models are general circulation models (GCMs) and Earth system models (ESMs). These models solve the discretised fundamental physical equations governing the behaviour of the atmosphere and ocean on a finite resolution grid. They also include representations of other components of the climate system, such as sea-ice processes and the carbon-cycle (in ESMs). A representation of the structure of a GCM is shown in figure 1.1. From an initial starting state, these models evolve the state of the climate system forward in time, outputting the future model state at specified times on a finite resolution model grid. GCM grid resolution is constrained by computational power and computer time afforded to the simulations, with typical horizontal resolution (spacing between grid points) of 100km ($\approx 1°$) in contemporary GCMs.

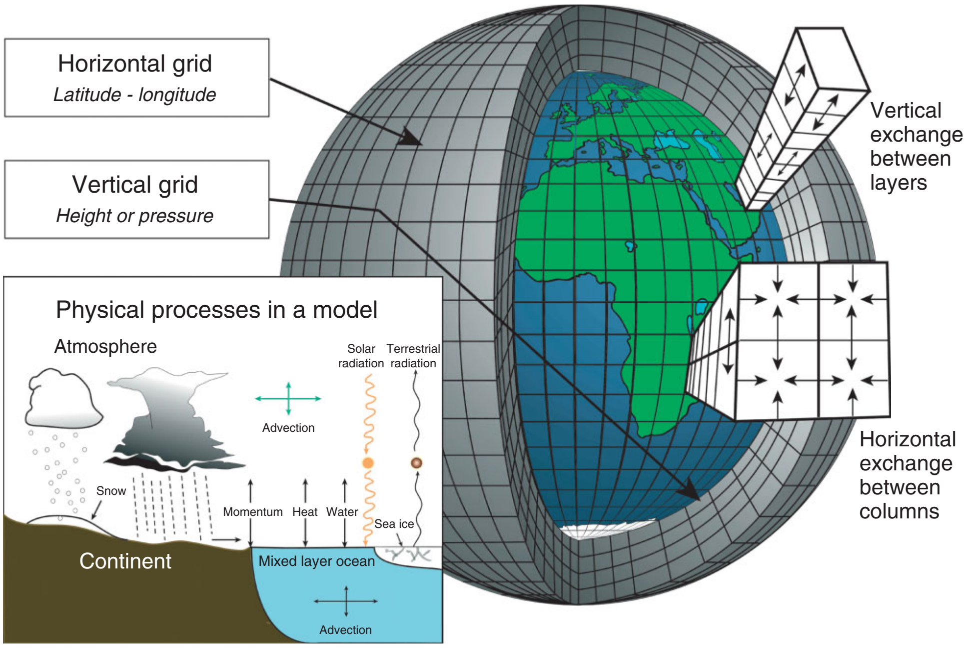


**Figure 1.1** – The structure of a typical GCM, denoting the 3 dimensional grid on which the large-scale variables are resolved. The sub-panel shows a representation of processes that occur on sub-gridscales and are therefore parametrised within the GCM. Reproduced from Edwards [2011], figure 2.

Current GCM grid resolution creates a number of challenges for simulation of the climate. Many processes which are important to the large-scale behaviour of the climate system occur on scales smaller than the spacing of the GCM grid. For instance, cloud formation, an important contributor to the overall albedo (reflectivity) of the planet, occurs on much smaller scales than GCM grid-boxes and is therefore not directly simulated by the solution of the fundamental fluid dynamics discretised on the GCM grid. To capture the behaviour of such processes it is necessary to parametrise their effect indirectly in terms of the variables that are resolved on the GCM grid. To do this, higher resolution models and observations are used, coupled with theory, to find relationships between the large-scale variables on the GCM grid and the behaviour of the small-scale processes.

Parametrisation introduces uncertainty into a GCM experiment. Firstly, there are multiple (arguably equally-justifiable) structural choices for the empirical parametrisation scheme used to relate the sub-gridscale process to the resolved variables. Secondly, these parametrisations require the introduction of several numerical constants, many of which are very difficult, or impossible, to constrain from observations or other process models. Uncertainty associated with the value of the sub-gridscale parameters is known as parametric uncertainty. Multiple parametrisation schemes are needed to capture all climatically important processes occurring on sub-gridscales. This can create complex interactions between parametric uncertainty in each parametrisation scheme, which propagates into uncertainty about the response of a GCM to an external forcing. The combination of these uncertainties in GCMs is known as *model uncertainty*.

**The multi-model ensemble (MME)**

Model uncertainty can be estimated by sampling different choices for parametrisation schemes of sub-gridscale processes. Comparing GCMs from different modelling centres around the world, run under the same forcing experiment, can give an indication of the impact of model uncertainty on the behaviour of the climate system (e.g Forster *et al.* [2013]). These models sample different choices for both the physical parametrisation schemes and for the dynamical core of the GCM (the numerical method for evolving the fundamental equations governing the fluid flow of the atmosphere and ocean forward in time). The group of these models is conventionally known as the *multi-model ensemble* (MME). However, as the MME is an *ensemble of opportunity* (not an ensemble intentionally designed to sample all possible modelling choices), the MME cannot be considered fully representative of model uncertainty. Indeed, as different GCMs have not been developed entirely independently of each other (and in some cases have shared code bases, shared personal and shared viewpoints on how to implement processes in a GCM) the *effective* number of independent models (in a statistical sense) within the MME is less than the number of GCMs [Knutti *et al.,* 2010]. In light of these facts the interpretation of the MME for probabilistic projection of climate change requires careful thought [Knutti, 2010].

#### Perturbed-parameter ensemble (PPE)

An alternative to exploring model uncertainty through the diverse but unstructured MME is to explore model uncertainty within a single GCM framework through a *perturbed-parameter ensemble* (PPE - also know as a *perturbed physics ensemble*). A PPE only explores a single parametrisation framework within a certain GCM but can sample many values for coefficients used within the model's parametrisation schemes. Despite being limited to a single GCM framework (mainly the UK Met Office's HadCM3 GCM), in many cases PPEs (all using the HadCM3 model) have been shown to span equal or greater uncertainty ranges than the MME [Murphy *et al.,* 2004; Stainforth *et al.,* 2005; Rowlands *et al.,* 2012]. An additional advantage of the PPE is that parameter space (and therefore model uncertainty) can be explored in a systematic and structured way [Sexton *et al.,* 2012], facilitating interpretation for probabilistic climate projection [Murphy *et al.,* 2009].

### 1.2.2 Models of reduced complexity

GCM simulations can be sufficiently complex that it can often be very difficult to deduce the cause of simulated phenomena within these comprehensive models. This creates a need for simpler, more idealised models of the climate system, that gain in tractability and comprehensibility what they lose in complexity and realism. These simple models can take a wide range of forms, from zero-dimensional energy-balance models of the entire climate system [Shine *et al.,* 2003] to 1 or 2 dimensional (typically vertical and meridional) models of the atmosphere that calculate the radiative-convective equilibrium of the atmospheric vertical column [Manabe & Strickler, 1964].

Additional advantages offered by a simple model formulation are the minimal computational resources required to conduct experiments with them. State-of-the-art GCMs can take many days to run, even on supercomputers, whilst in contrast integrations with simpler models of reduced complexity can be run in a matter of seconds on desktop processors. Differences in future climates are highly dependent on *scenario uncertainty*, which is associated with the many differing possible future pathways of anthropogenic emissions of climate forcing agents (e.g. see Van Vuuren [2011]). Exploring the climatic response to this uncertainty (which becomes the dominant uncertainty on century long timescales [Hawkins & Sutton, 2009]) in sufficient detail to be able to conduct policy analysis prohibits the use of computationally expensive GCMs and therefore requires reliable results from simpler representations of the climate system.

## 1.3 The importance of uncertainty in the climate response to $CO_2$

This section introduces one of the primary and most directly policy-relevant challenges in climate science, namely uncertainty in the response of the climate system to enhanced concentrations of $CO_2$, which is to be the topic of the rest of this thesis. Whilst $CO_2$ is the largest radiative forcing on the climate system (estimated to contribute 1.82Wm$^{-2}$ in 2011 [Myhre, 2013a]), $CO_2$ is not the only anthropogenic forcing inducing climatic changes. Total radiative forcing in 2011 is estimated to be 2.29Wm$^{-2}$ when the forcing of $CO_2$ is combined with additional positive radiative forcings (such as from other well-mixed greenhouse gases) and incorporating the negative radiative forcing of aerosol compounds (-0.9Wm$^{-2}$ in 2011). Whilst there remains substantial uncertainty in the present-day forcing of aerosols (consensus estimates range from -1.9Wm$^{-2}$ to -0.1Wm$^{-2}$ [Myhre, 2013a] largely due to uncertain knock-on consequences of aerosol emissions for the radiative effect of clouds) and their future evolution [Meinshausen *et al.,* 2011; Shindell *et al.,* 2013], unlike $CO_2$ these forcing are short-lived in the atmosphere and are therefore less important than the response to $CO_2$ for long-term and permanent climate change [Smith *et al.,* 2012]. I therefore choose to focus the work in this thesis on understanding the climate response to $CO_2$ as the primary goal.

Characterising uncertainty about the climate system's response to an emissions scenario is an essential aid to making informed policy responses to climate change. As will be reviewed in chapter 2, there are many interacting physical mechanisms that determine the response of the climate system to $CO_2$. These responses are difficult to constrain from either observations of the climate system or simple and complex climate models. Charney *et al.* [1979] offered a first consensus estimate of a common measure of the climate response to $CO_2$, the *equilibrium climate sensitivity* (ECS - the equilibrium GMST change from a doubling of $CO_2$ concentrations, see section 2.4), to lie between 1.5K and 4.5K, a broad range of uncertainty. More than 30 years later, despite large advancements in both computational power and understanding of climate system processes, the latest assessment report of the Intergovernmental Panel on Climate Change (IPCC) [IPCC, 2013] assessed the *likely* ECS range (>66% belief that the true value of the parameter lies within the given range) as the exact same interval (although Charney *et al.* [1979] didn't assign a probability to this interval), demonstrating the fundamental difficulties in constraining the magnitude of the climate response to $CO_2$.

Despite broad ranges of uncertainty in the climate response, this uncertainty is not, in itself, a reason to avoid taking any action to mitigate anthropogenic climate change, as formalised Bayesian risk analysis frameworks exist to analyse decision-making under physical climate response uncertainty (e.g. Tebaldi *et al.* [2005]). Indeed, uncertainty about the physical response of the climate system to forcing needs to be incorporated into evaluations of mitigation policies seeking to limit realised changes in the physical climate system. Mitigation policy frameworks that are fundamentally resilient to physical climate response uncertainty can help offset the knock-on impact of this uncertainty on realised changes in the climate system (see chapter 7).

Within a risk analysis framework, it is essential to understand and capture the full uncertainty in important quantities as completely as possible to enable optimal decision-making. Specific to the climate change case, the shape of the global economic damage function associated with warming (which is likely upwardly concave with increasing temperature) gives particular importance to the quantification of the upper tail of the probability distribution for climate response [Weitzman, 2011]. The non-linearity in temperature of climate change damage curves arises due to, amongst other things, damages associated with changes in extreme event frequency, as each subsequent unit increase in the time-mean of geophysical distributions gives rise to a greater increase in the probability of extreme events than the previous one. The broadening or narrowing of the wide uncertainties in the mean climate response can therefore have a large impact on the assessment of expected damage under a specific emissions scenario.

As a specific example of the importance of uncertainty quantification in the climate response to $CO_2$, figure 1.2 shows probability density functions (pdfs) for the transient climate response (TCR - see section 2.4). The TCR (GMST change at the time of doubling of $CO_2$ concentrations under a 1% annual increase) is the measure of physical climate response most relevant to peak temperatures under future $CO_2$ emissions scenarios in which emissions peak and eventually decline (assuming that significant non-linearities do not exist in the Earth system at temperatures reached under peak and decline scenarios). Assuming, for simplicity, that the uncertainty in the thermal response to radiative forcing dominates uncertainty in the response of the carbon-cycle to anthropogenic carbon emissions [Gillett *et al.,* 2013], uncertainty in the instantaneous human-induced temperature response to the cumulative

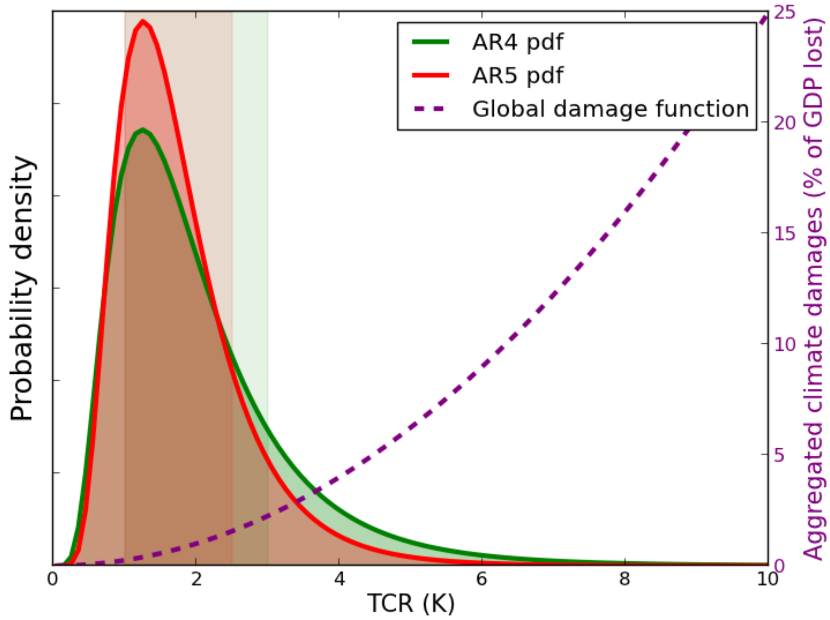


**Figure 1.2** – Log-normal pdfs for the transient climate response (TCR). The green/red curve shows the pdf consistent with the inferred *likely* uncertainty ranges (vertical shaded bands) for the TCR in the 4th/5th IPCC Assessment Report respectively. The purple curve (right hand axis) shows a representative quadratic damage function taken from Nordhaus [2008].

emissions required to double the pre-industrial atmospheric $CO_2$ concentration is determined by the uncertainty in TCR. Therefore, a pdf in TCR could be assumed to represent a pdf in the peak warming realised under the emissions scenario. Figure 1.2 shows log-normally distributed pdfs (a typical shape for TCR pdfs due to the existence of low probability, high response tails in a positive definite variable [Pueyo, 2012]), that have been explicitly chosen to have *likely* probability intervals consistent with the *likely* intervals given for the parameter in the IPCC 4th (AR4) and 5th (AR5) Assessment Reports respectively [IPCC, 2007; IPCC, 2013] (the AR4 interval is inferred from the *likely above* and *likely below* statements about TCR; see chapter 4).

Using these pdfs it is possible to calculate the difference in the *expected* damage (at time of doubling, not the net present value of all future damages discounted in time) due to climate change (assuming a specific known form for the damage function) under both of these two different distributions. Here, a quadratic globally-aggregated damage curve is taken from Nordhaus [2008]. Using the aggregated measure of *likely* intervals, the only difference between the two log-normal pdfs is the location of the upper bound of the *likely* probability interval (2.5 as opposed to 3K), whilst the lower *likely* interval bounds are identical and the most likely TCR values in the two distributions are very similar (both approximately 1.25K). The normalisation of the probability distribution means that the AR4 pdf has more probability in the high response tail than the AR5 distribution. When evaluating the expected value of the damage curve under these TCR probability distributions, the additional probability in the tail for the AR4 pdf has a substantial impact on the result. The expected damage under the AR4 distribution is 14.6% of global GDP, however, under the AR5 distribution the expected damage shrinks to 9.9%, a fractional change of nearly a third. Such a change in the expected damage could have a substantial knock-on effect on further climate policy and cost-benefit analysis that might be conducted with the result and could substantially alter optimal mitigation and adaptation decisions.

The idealised example shown above clearly demonstrates the importance for decision making of constraining uncertainty in physical climate response as completely as possible. As shown, two different pdfs that have very similar best-estimates of the physical climate response could potentially lead to very different decision-making implications under a risk framework, due to the sensitivity to the extremes of the probability distribution described above. This example illustrates the motivation for developing physical understanding of the climate system's response to external forcing and probing the limits of this uncertainty in order to quantitatively describe the entire pdf of the response as fully and accurately as possible, as will be addressed in the subsequent chapters of this thesis.

Whilst the above example focused on the impact of constraining uncertainties at the high response end of the pdf, understanding the lower limit of the climate response uncertainty (one of the main topics of the later chapters of this thesis) is also important for decision making. Estimates of minimal costs of climate damages under different emissions scenarios as well as alleviation of risks of potentially costly over-mitigation (where profitable energy infrastructure and resources are written-off based on a overly-high estimate of the climate response) are highly dependent on a correct characterisation of the lower end of the climate response uncertainty distribution. A recent study suggested that certain climate impacts may be sigmoidal (saturating) with temperature [Ricke *et al.,* 2016], leading to the existence of both a low warming/high mitigation and high warming/low mitigation equilibrium between marginal mitigation and adaptation costs. In such a case, fully and exhaustively constraining the lower limits of climate response will be essential for assessing the feasibility of the low warming/high mitigation equilibrium, as it also will be for assessing efficient mitigation pathways that could be compliant with the 1.5K peak warming ambition of the international climate policy process [UNFCCC, 2015].

## 1.4 Outline of the thesis

This thesis presents a physically-based investigation into uncertainty in the climate response to $CO_2$, using observations, simple energy-balance models and GCM PPE experiments. The substantial mitigation policy implications of constraining physical response uncertainty, discussed in this chapter, serves as the underlying motivation for this work. This section describes a brief roadmap for what is to follow in the remaining chapters of the thesis.

Chapter 2 conducts a discussion and literature review of the major uncertainties in the physical climate response, in order to provide the necessary background to the mechanisms and methods relevant to the experiments conducted in this thesis. In particular, this chapter highlights the uncertain response of low cloud as the major determiner of uncertainty in the global climate response to $CO_2$ and is a key physical driver of the results shown later in chapters 5 and 6.

Chapter 3 describes the methods of perturbed physics ensembles that are used to investigate climate response uncertainty in parts of this thesis. This chapter details the parametrisation schemes perturbed in the HadCM3 PPE, before going on to describe features of the climate response in an initial ensemble of PPE experiments that aim to sample a broad range of the uncertainties in the components of climate feedback and forcing. PPE experiments are conducted for both instantaneous doubling and 1%/yr concentration increase. A long tail of high climate sensitivity is found in this PPE (with estimated equilibrium warmings greater than 9K), but an abrupt end to the distribution around 2K is seen on the lower end of the range.

Chapter 4 examines estimates of climate response derived from observations using simple energy-balance models of the climate system. These models are used to assess which metric of climate response is most important for future climate projections. The chapter then estimates the climate response from observations of the 2000-2009 planetary energy budget and shows how these estimates are sensitive to the structure of the simple models used. Finally, the chapter introduces the concept of *realised warming fraction* (RWF) and highlights an area of climate response space that is consistent with observationally-based energy-balance estimates but not sampled by GCMs. The majority of the work presented in this chapter has been published in *Climatic Change*.

Based on the discrepancy between observationally-based and GCM-simulated ranges of climate response that was highlighted in chapter 4, chapter 5 describes a statistical methodology to emulate the response of the HadCM3 PPE and an attempt to create an ensemble of PPE members that have ECS of less than 2K. The properties of the climate response of this PPE, and in particular the set of GCMs that are simulated to have ECS<2K, are analysed and the implications for the plausibility of ECS<2K discussed. GCMs are found that have reduced skill from the standard HadCM3, but would not be complete outliers from the CMIP5 distribution on many common metrics.

Chapter 6 goes on to describe the physical mechanisms underlying the differences in climate response shown in the ensemble of chapter 5, using a sub-set of these GCMs indicative of the range of behaviour across the ensemble. Cloud feedbacks in these GCMs are decomposed into different physical regimes using a stability index, and the physical mechanisms underlying the response of tropical low clouds (the greatest source of difference in climate response across the predicted low ECS ensemble) are analysed in detail. The quality of the control climate simulation in these models is also analysed. A combination of the work in chapters 5 and 6 is being prepared for submission to *Climate Dynamics*.

This thesis shows robust and systemic uncertainty in the response of the climate system to increases in atmospheric $CO_2$, consistent with the wider literature. Chapter 7 describes a new adaptive mitigation policy framing that can be constructed to be resilient to uncertainty in the physical climate system response. This policy advocates a focus on control variables that are capable of limiting the cumulative stock of $CO_2$ emissions into the atmosphere and indexing evolution of these variables to an updated index of the emerging climate response. This chapter also introduces a new simple climate-carbon-cycle model that is used to directly connect emissions of $CO_2$ and GMST changes and captures the physical feedbacks between the climate and carbon-cycle apparent in ESMs. Work from this chapter is under review at *Atmospheric Chemistry and Physics* and has also been submitted to *Environmental Research Letters*.

Finally, this thesis finishes with a short conclusions chapter, chapter 8, summarising the previous chapters and highlighting some potential areas for future work.

# CHAPTER 2

# Climate forcing and feedbacks

## 2.1 The forcing-feedback framework

On the planetary scale, the response of the climate system to external forcing agents can be analysed in terms of changes in energy flows within the global climate system. This long-standing and well-established framework has been useful for understanding the sensitivity of the climate system to the radiative forcing of increased $CO_2$ concentrations.

The global-mean perturbation energy-balance of the climate system is canonically summarised with the zero-dimensional equation,

$$N = \Delta F - \lambda \Delta T, \tag{2.1.1}$$

where $N$ is the net energy imbalance of the climate system (equivalent to the heat uptake by the oceans on climatological timescales due to the much larger heat capacity of the oceans than the land and atmosphere), $\Delta F$ is the *radiative forcing* of the climate system (defined as the change in net top-of-atmosphere (TOA) energy-balance that occurs when an instantaneous radiative perturbation is applied to the climate system after allowing for instantaneous adjustments) and $\lambda \Delta T$ is the radiative response of the climate system associated with changes in the GMST of the climate system, $\Delta T$ [Gregory *et al.,* 2004]. $\lambda$, known as the *climate feedback parameter* is a measure of how strongly changes in GMST, responding to external forcing, are associated with the radiative response of the climate system. Its value (determined by climate system responses to increased $\Delta T$ that impact the global energy budget) determines how much GMST warming the climate system will undergo to completely respond to the external forcing and reach a new equilibrium state.

The climate feedback parameter can be decomposed into multiple components associated with the radiative impact of temperature-driven changes in climatic variables,

$$\Delta R = \lambda \Delta T = \Delta T \left( \frac{\partial R}{\partial T} + \sum_i \frac{\partial R}{\partial X_i} \frac{\mathrm{d}X_i}{\mathrm{d}T} \right), \tag{2.1.2}$$

where $\Delta R$ is the total global-mean radiative response of the climate system to the forcing and $X_i$ are global-means of other radiatively important variables, such as planetary albedo or water vapour concentrations [Hansen, 1984]. All terms in equations 2.1.1 and 2.1.2 can be considered as functions of time and evaluated either for transient states or for the ultimate equilibrium response to a perturbation. Section 3.4.2 in Pierrehumbert [2010] offers a clear disucssion of the forcing-feedback framework for further consultation.

For the climate system to be stable, the net radiative response to forcing must oppose the forcing, therefore, as defined in equation 2.1.1, $\lambda$ must be positive (a net negative feedback). The first term within the parentheses on the right hand side of equation 2.1.2, $\frac{\partial R}{\partial T}$, is known as the Planck feedback and is the change in the outgoing infra-red radiation to space, assuming a uniform temperature change throughout the atmosphere with all other radiatively-important climatological variables remaining unchanged. From the Stephan-Boltzmann law it is clear that a black body emits more radiation as it warms and therefore the term $\frac{\partial R}{\partial T}$ is positive (negative feedback), with a value of about 3.2Wm$^{-2}$K$^{-1}$ as simulated in the latest GCMs [Vial *et al.,* 2013]. The deviation of the total feedback parameter away from the Planck feedback is driven by additional negative and positive feedbacks associated with changes in other radiatively-important climatological variables. The time evolution of these other components of the radiative feedback will determine the magnitude of the time-dependent response of the GMST to an effective radiative forcing.

There are multiple assumptions made in employing this forcing-feedback framework as formulated in equation 2.1.2. These assumptions are detailed in Sherwood *et al.* [2014a]. The main one being: an assumption of linear global responses to perturbations, such that the partial derivatives in equation 2.1.2, can be evaluated as time-independent constants, and the additional assumption that different patterns of perturbation resulting in identical global-mean changes have identical radiative effects on the system. These assumptions require radiative responses to changes in climate system properties to be independent of the background climate state, an assumption which must break down in climates substantially changed from the present-day. For instance, planetary albedo must remain greater than or equal to zero and therefore a linear relationship between albedo change and warming must eventually break down at low albedo. Non-Planck feedback mechanisms can lead to the total radiative response of the climate system depending on the pattern of warming induced by the perturbation, as many radiatively-important climatic variables experience spatially inhomogeneous changes dependent on the spatial pattern of temperature change [Hansen *et al.,* 1997].

## 2.2 Effective radiative forcing, time dependent climate feedbacks and efficacy of ocean heat uptake

The concept of *instantaneous radiative forcing* is an idealised concept in which instantaneous perturbations are made to a model of the climate system (normally the introduction of a radiative heating perturbation altering energy flows within the atmospheric column). The *instantaneous radiative forcing* is then defined as the change in net radiative balance at some specified height in the vertical profile of the atmosphere (normally either the TOA or the tropopause) upon introduction of the perturbation, without letting the climate system respond to the perturbation. It is important to note that the concept of an instantaneous perturbation without response is not directly observable in the real world, as the climate system begins to respond instantaneously to any perturbation.

Part of the motivation for developing the concept of radiative forcing is the desire to be able to develop a common metric in which the climatic perturbing potential of different radiative important forcing agents can be compared. Specifically, in light of the policy challenge of contemporary climate change, the desire for a metric that accurately encodes the potential for change in the GMST in response to the forcing has been an underlying driver of the development of the radiative forcing concept. However, for some forcings there are *rapid adjustments* (changes that occur on timescales that are much shorter than timescales of climate response and are direct consequences of the introduction of the perturbation) that very quickly modify the radiative imbalance associated with the perturbation. Determining the radiative forcing after allowing for these rapid adjustments to take place gives radiative forcing values that are much better measures of the ultimate potential for the perturbation to cause changes in the surface climate system. Accounting for these rapid-adjustments has led sequentially to the development of *strato-*

*spheric adjusted radiative forcing* and *effective radiative forcing*, calculated after allowing stratospheric temperatures to adjust to the perturbation (stratospherically adjusted radiative forcing) and additionally allowing for fast adjustments in the troposphere, such as direct adjustments to cloud fields, associated with the forcing to occur (effective radiative forcing) [Hansen *et al.,* 2005; Myhre, 2013a].

In recent years, the concepts of fast adjustments and the possibility of time varying climate feedbacks have often been conflated, leading to some confusion in the literature. In particular, confusion arises as methods of calculating forcings and feedbacks through regressions of the TOA energetic imbalance against GMST in climate simulations [Gregory *et al.,* 2004] can lead to different values for the effective radiative forcing and linear climate feedback parameter depending on what time period they are calculated over. This is partly a consequence of the fact that there is no clear distinction between the timescales of rapid adjustments and the timescales of the fastest climate feedbacks (e.g. Colman & McAvaney [2011]). Actual time-dependence of climate feedbacks (for example from temporal and hemispheric variations in cloud feedback [Senior & Mitchell, 2000; Gregory *et al.,* 2004]) has also been documented in the literature. This leads to a curved relationship between the climate's energetic imbalance and the GMST response, different from the canonical linear relationship. This is shown pictorially in figure 2.1.

Additionally, Winton *et al.* [2010] has shown that variations of the *efficacy* of the ocean heat uptake between GCMs, when the ocean heat uptake is thought of as an additional radiative forcing on the surface energy-balance (non-unity efficacy derives from the fact that differing patterns of ocean heat uptake with equal global magnitude will result in different radiative feedbacks in the global-mean, equivalently to as with forcings), can be used to explain the apparent time-dependence of the climate feedbacks without the need to actually invoke time-dependent global feedback parameters. Winton *et al.* [2010] state the large majority of the time-dependence of the climate feedback in global climate models can be explained by a time-invariant ocean heat uptake efficacy factor. However, as patterns of ocean

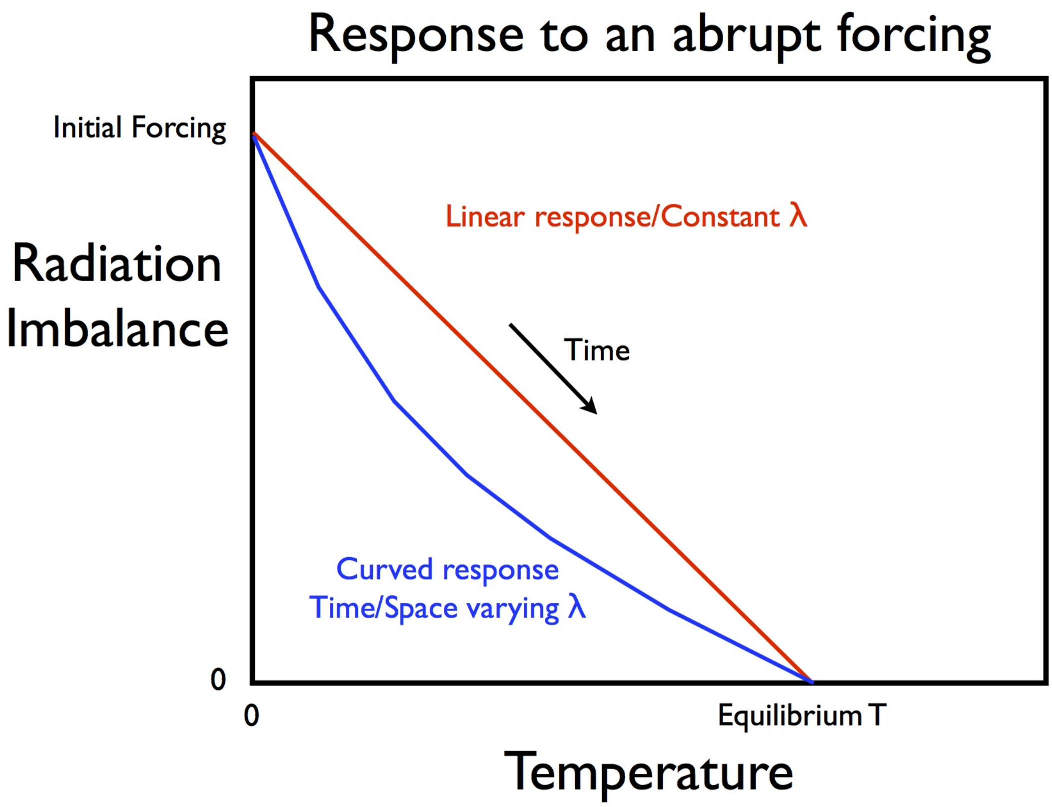


**Figure 2.1** – A pictorial representation of the relationship between the global energetic imbalance of the climate system (y-axis) and the GMST change (x-axis) with linear (red line) and non-linear (blue line) feedbacks. Reproduced from Real Climate, originally from Jonathan Gregory.

heat uptake change in time, associated with changes in the oceanic circulation, the efficacy of the ocean heat uptake also varies temporally, even with time invariant radiative feedbacks.

## 2.3 Uncertainty in components of climate feedback

Conventionally, the climate feedback parameter is decomposed into several different components,

$$\lambda = \sum_i \lambda_i = \lambda_p + \lambda_{wv} + \lambda_{lr} + \lambda_{albedo} + \lambda_{cl}, \tag{2.3.1}$$

where $\lambda_p$ is the Planck feedback (assuming that temperatures are changing uniformly throughout the troposphere), $\lambda_{wv}$ the feedback associated with the atmospheric water vapour response to warming, $\lambda_{lr}$ the feedback associated with changes in the vertical temperature lapse-rate of the atmosphere in response to warming, $\lambda_{albedo}$ the feedback associated with changes in the surface albedo in response to warming and $\lambda_{cl}$ the feedback associated with changes to clouds in response to warming. The real climate system contains other feedback processes, including ones acting over longer timescales to those detailed in equation 2.3.1. These include carbon-cycle feedbacks in response to warming such as a change in the ability of the biosphere to take up carbon with increased temperature [Gregory *et al.,* 2009]. Conventionally, due to the historical development of climate system modelling, the physical climate system has, until recently, been modelled and analysed with concentrations of radiatively active atmospheric gases prescribed. Therefore, these carbon-cycle feedbacks are not included in the estimates of the physical climate feedbacks and the metrics that are typically used to summarise the climate response. Further discussion of relevant carbon-cycle feedbacks is left until chapter 7.

Understanding the magnitudes and uncertainties associated with these different components of climate response has been an essential goal of intercomparison projects between different GCMs from around the world (e.g. the CMIP3 [Meehl *et al.,* 2007] and the CMIP5 [Taylor *et al.,* 2012] intercomparison projects). Figure 2.2 reflects ranges of components of the feedback parameter, described in equation 2.3.1, in recent versions of GCMs as summarised in the 5th Assessment Report of the IPCC [Flato, 2013]. Understanding and constraining the physics of each of these feedback components is vital to improving understanding of the climate system's response to perturbations. In this section I briefly examine the current state of knowledge underlying each of these feedback components in turn, including the dominating uncertainty in global cloud feedbacks.

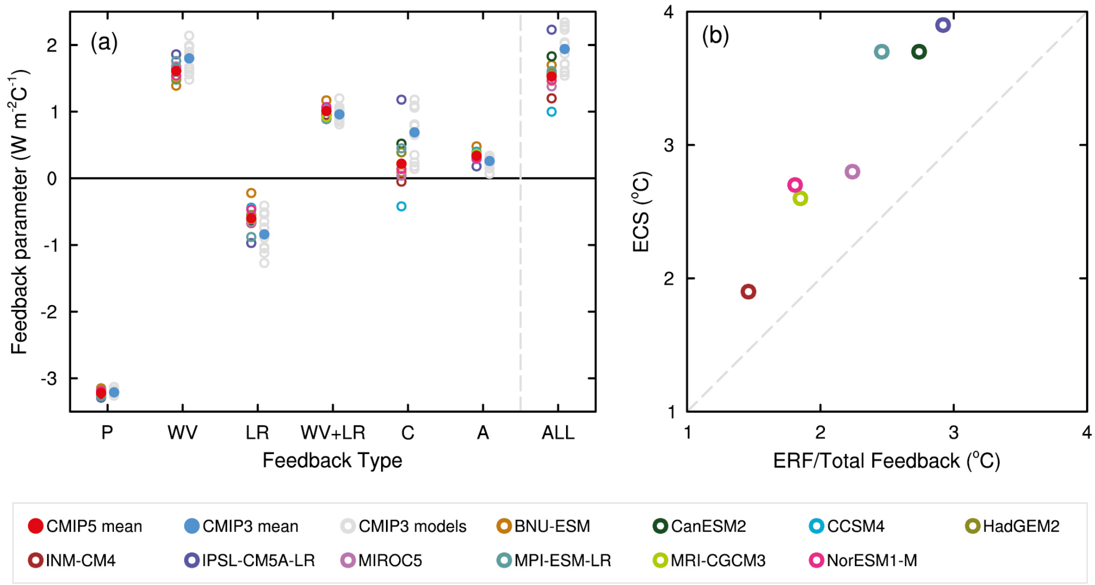


**Figure 2.2** – Breakdown of the feedback parameter into different components in CMIP3 and CMIP5 GCMs. This image is taken from figure 9.43 in Chapter 9 of IPCC AR5 WG1 [Flato, 2013]. The water vapour and lapse-rate feedbacks are combined in column (WV+LR) and the ALL column represents the sum of all feedback components except for the Planck feedback. The sign convention used is opposite to that used in equation 2.1.1 and therefore a positive feedback parameter component represents a positive feedback on the system. The right hand panel shows how a linear estimation of the effective radiative forcing (ERF) and total feedback parameter can lead to an underestimation of the ECS, in part driven by the time-dependence of the feedback components.

### 2.3.1 The water vapour and lapse-rate feedback

It has long been known that GCM simulations of a warming climate show a roughly unchanged distribution of relative humidity [Manabe & Wetherald, 1967; Soden & Held, 2006]. A possible explanation for this feature of climate simulations is that relative humidity is set by the point of last saturation of a moist parcel of air moving with the atmospheric circulation [Sherwood *et al.*, 2010]. Water vapour is lost from an air parcel when atmospheric conditions are encountered on its trajectory that have a saturation vapour pressure lower than the vapour pressure of the air parcel. Hence, the specific humidity at any subsequent point is set by the specific humidity at the last point of saturation encountered on its trajectory [Pierrehumbert *et al.*, 2007]. Assuming that the atmospheric circulation does not change substantially under a warming climate, the specific humidity at any point in space will evolve with changes to the saturation specific humidity at the point of last saturation on the mean trajectory of air parcels arriving at that point. Therefore, as saturation vapour pressure increase with temperature according to the Clausius-Clapeyron relationship, the atmospheric specific humidity also increases under the Clausius-Clapeyron relationship. If, at any location in the atmosphere, the change of local saturation vapour pressure is not substantially different from the change at the point of last saturation for air parcels arriving at this point, the local relative humidity will remain more or less constant under a warming climate [Sherwood *et al.,* 2010]. Water vapour has a very strong greenhouse effect and hence the water vapour feedback is a substantial amplifier of warming under climate perturbations. The magnitude is relatively robust and well constrained in models and can be understood in terms of fundamental physical considerations [Ingram, 2010].

The atmospheric lapse-rate (the rate of decrease in atmospheric temperature with altitude) is, in the tropics at least, expected to change as a moist adiabatic lapse-rate evolves under warming, due to convective motions that act to restore atmospheric vertical profiles to the moist adiabatic lapse-rate in the tropical regions. The Planck feedback assumes that the troposphere warms uniformly, however, the evolution of the moist adiabat gives a inhomogeneity to the vertical structure of warming in the atmosphere. This is important for climate feedbacks, as changes in the lapse-rate alter how much radiative response (at the TOA) is generated per unit of GMST change. A flattening of the lapse-rate would increase the radiative cooling to space for each unit of GMST change, whereas a steepening would mean that more GMST change would be required to provide the equilibrium radiative response to the initial radiative perturbation.

Lapse-rate and water vapour feedbacks tend to be anti-correlated in contemporary GCMs (as shown by the reduced uncertainty in their combined sum in figure 2.2). This is physically well understood as the two feedbacks are dependent on each other and the relative humidity of the atmosphere [Hansen, 1984; Ingram, 2013]. Simply summarised, increased water vapour content raises the altitude (lowers the temperature) from which infra-red radiation escapes to space. Whilst a steepening of the lapse-rate would lower the temperature at the same level, the change in lapse-rate would also reduce the specific humidity at that altitude, reducing the infra-red opacity of the atmosphere and therefore lowering the altitude (and increasing the temperature) from which the infra-red radiation escapes to space, a compensating effect [Ingram, 2013].

### 2.3.2 The surface albedo feedback

As the climate warms, highly reflective snow and ice begins to melt, revealing lower albedo surface beneath, allowing more solar radiation to be absorbed by the climate system. This effect reduces the global mean albedo of the planet and acts as a short-wave clear-sky positive feedback on the planet. The magnitude of the surface albedo feedback is dependent on the complex interaction between warming and oceanic currents that drive reductions or increases in sea-ice cover in both hemispheres and changes to their spatial and seasonal distributions.

### 2.3.3 Cloud feedbacks

Clouds, which form when atmospheric air becomes either moist or cold enough to allow sufficient super-saturation for vapour condensation onto cloud condensation nuclei, have a substantial additional effect on radiative transfer in the atmosphere over and above gaseous water vapour.

Cloud processes vary across a vast range of scales in the climate system, with important cloud-relevant processes such as organised convection occurring on spatial scales as small as tens of kilometres, much smaller than the grid resolution in current GCMs. Additionally, cloud properties are the result of interactions between microphysical, dynamical and thermodynamical processes and hence are inevitably highly complex phenomena in GCMs [Stephens, 2005]. The accurate simulation

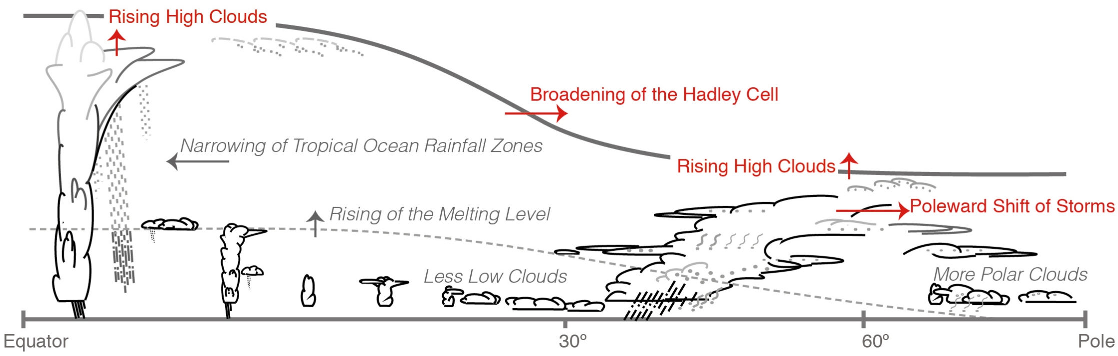


**Figure 2.3** – Robustly simulated cloud responses in GCMs. Taken from figure 7.11 in chapter 7 of IPCC AR5 WG1 [Boucher, 2013]. The thick grey line represents the tropopause and the thin grey line the melting level. Mechanisms in red are robust positive feedback mechanisms and shown in grey are mechanisms that have little effect or for which the sign of the feedback is uncertain.

of cloud formation, their radiative effects on today's climate and cloud changes in response to a climate forcing are all very difficult challenges for GCMs. It is therefore unsurprising that the magnitude, and even the sign, of the net cloud feedback is uncertain in figure 2.2 and that cloud feedbacks are the dominant uncertainty in the overall feedback parameter. It has also been shown that cloud feedbacks are a large contributor to the overall model spread of transient climate response to ramped forcing [Dufresne & Bony, 2008].

There are a host of different cloud feedback mechanisms that are both hypothesised and/or simulated by GCMs. High clouds are expected to rise under a warming of the troposphere but the temperature at the cloud top is expected to remain roughly constant [Hartmann & Larson, 2002; Zelinka & Hartmann, 2010]. As the surface beneath the cloud warms without an increase in the temperature at which infrared energy is lost to space, this creates a positive long-wave feedback. GCMs also simulate a reduction in high and mid level clouds in equatorial and sub-tropical regions and a poleward shift in the locations of clouds, at least in part associated with changes in large scale dynamics of the atmospheric circulation (such as a broadening of the Hadley Cell) [Wetherald & Manabe, 1980].

**Low cloud feedbacks**

Low clouds (typically defined as clouds that have cloud tops at altitudes below 680hPa [Webb *et al.,* 2013]) are typically optically thick and hence have a negative short-wave cloud radiative effect (associated with the reflection of solar radiation to space) whilst having a small greenhouse effect on the surface climate system (as they are relatively close to the temperature of the surface and hence do not have a strong effect on the temperature at which long-wave radiation escapes to space).

Feedbacks involving low cloud have been consistently shown to be the most important determiner of the overall cloud feedback parameter and hence the overall climate sensitivity [Randell, 2007; Boucher, 2013; Webb *et al.,* 2013]. Most GCMs simulate a decrease in low cloud cover as a response to GMST warming, but the magnitude of this decrease differs between GCMs [Stephens, 2005; Vial *et al.,* 2013]. In particular, feedbacks involving marine boundary layer clouds in low-latitude subsidence regimes are a large source of uncertainty in global cloud feedbacks [Bony & Dufresne, 2005].

Feedbacks of the low cloud subsidence regime have often been investigated using idealised representations of the tropical climate. Figure 2.4, taken from Bony *et al.* [2006] shows a typical representation of the tropical convective regions and the subtropical subsidence regions, along with a two-box model that has often been used to study the dynamics of cloud feedbacks through interactions between the the warm and cool pools of the tropics [Pierrehumbert, 1995; Miller, 1997; Larson *et al.,* 1999]. The height of cloud tops is empirically related to the stability of the lower troposphere through various different metrics [Miller, 1997; Wood & Bretherton, 2006; Bony *et al.,* 2006]. Locally, the convective potential of the atmospheric column above a region of tropical ocean is dependent on the sea-surface temperature. However, when examining low stratocumulus clouds (prevalent over large regions of subsidence on the eastern side of low-latitude ocean basins, where subsidence, and therefore heating, of the lower atmosphere is the greatest, strengthening the inver-

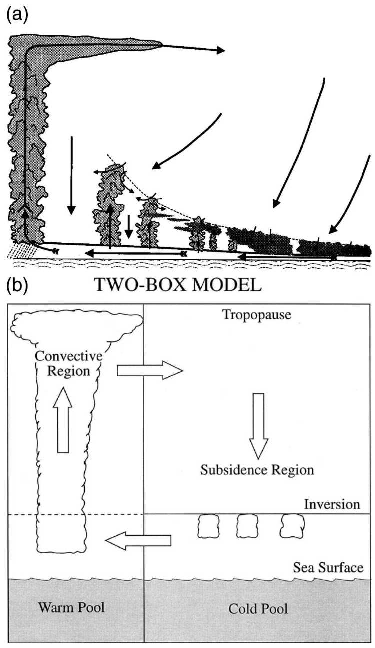


**Figure 2.4** – Figure 3 from Bony *et al.* [2006]. Panel a) shows the conventional representation of the tropical circulation regime. The dotted line represents the cloud top level associated with the sub-tropical temperature inversion, caused by a heating of the atmosphere from above by large-scale subsidence and a cooling from below due to boundary layer mixing. Panel b) shows an idealised model, that decomposes this circulation into two boxes, allowing the cloud response under warming to be analysed as a function of the heat and water vapour budgets in the two boxes and the associated transports between them.

sion), the cloud fraction is also dependent on the temperature of the atmosphere above the inversion level, which is nearly uniform across the sub-tropical region due to the effects of the large-scale circulation [Klein & Hartmann, 1993; Miller, 1997]. The potentially differing responses of the surface temperature, boundary layer and free-troposphere to $CO_2$ radiative forcing can provide competing feedback mechanisms on tropical low clouds. Differing simulations of the sign and relative strength of these feedback mechanisms account for the uncertainty in the sign of net tropical low cloud feedback in the latest GCMs. This competition of mechanisms may make the analysis of low cloud feedbacks dependent on the choice of stability metric, as different metrics may reproduce the present day variations in low cloud cover equally well, but may differ in their prediction of the response of low clouds under global warming [Wood & Bretherton, 2006]. GCMs often have limited vertical resolution near the boundary layer and crude representations of discrete cloud levels may fail to resolve important processes, such as shallow convective and turbulent motions, that contribute to the representation of boundary layer clouds despite the best efforts of GCM parametrisation schemes.

Unlike mid and high altitude cloud feedbacks (discussed previously) there is in general no clearly dominant feedback mechanism underlying low cloud response across GCMs. There are however, a number of different proposed physical mechanisms that may operate to constrain the low cloud feedback individually or in tandem. One proposed mechanism is that under a warmer climate the evaporation increase from the ocean (driven by changes in the air-sea temperature difference, ocean winds and near-surface humidity) will be insufficient to compensate for the increases in dry air entrainment into the marine boundary layer in a warmer climate, leading to an overall reduction of boundary layer moisture availability and hence low cloud cover in subsidence regions [Webb & Lock, 2013]. Sherwood *et al.* [2014b] also hypothesised that the rate of change of boundary layer dry air entrainment (through mid to low troposphere mixing) in CMIP5 models was associated with the magnitude of this mixing process in the present day climate, and therefore proposed an emergent constraint from observations on climate sensitivity towards the higher end of the currently accepted ranges.

In summary, low cloud feedbacks are particularly difficult to constrain due to the multiple driving processes and the challenge of accurately representing the local effects of these processes in typical GCM girdboxes. Due to these challenges, low cloud feedbacks are likely to be particularly sensitive to perturbations in subgridscale parameters and differences in parametrisation schemes across GCMs.

**Relation to observational constraints**

Observational constraints on the cloud feedback are difficult to derive and come with large uncertainties associated with the difficulty of identifying the signal of forced cloud changes relative to the fractionally large noise from natural variability. Another particular challenge of attempting to constrain cloud feedbacks from the observational record is the need to assume the relevance of cloud changes on inter-annual time cycles to changes under $CO_2$-driven warming. The use of inter-annual cloud changes has been shown to generally not be of use in estimating cloud feedbacks in GCMs [Dessler, 2010]. GCM-based evidence also suggests that cloud feedbacks in response to non-$CO_2$ forcing agents (at least for the volcanic forcing) are different to the feedbacks in response to $CO_2$ [Yokohata *et al.,* 2005], creating further complications in attempts to derive $CO_2$-driven cloud feedbacks from the observational record, where a variety of different forcing agents have been active in driving the observed change.

Turning to model-based observational constraints, GCMs do show relationships between observable elements of present day cloud properties and $CO_2$-induced feedback mechanisms [Yokohata *et al.,* 2010; Tett *et al.,* 2013]. Whilst the validity of these constraints have not been well-established in the literature [Klein & Hall, 2015] and such relationships could be specific to individual GCMs, the prospect of deriving an emergent constraint between aspects of cloud feedback and present day cloud properties offers what could be a fruitful way forward towards further constraining the total cloud and low cloud feedbacks.

## 2.4 Summarising the climate response to forcing

Within the framework described above, the forcing of $CO_2$ and the resulting climate feedbacks are typically combined in two standard metrics. The need for the physical climate response to be so summarised in so few metrics comes in part from the need to reduce complexities of forcing and feedback down to aggregated metrics that are easily communicable to policy makers and other non-technical parties. Many simple, but often-used, policy analyses of anthropogenic climate change incorporate the physical climate response solely through use of a single number summarising the sensitivity of the climate system to elevated concentrations of $CO_2$.

The *equilibrium climate sensitivity* (ECS) is a measure of the ultimate GMST response to a radiative perturbation - normally implicitly assumed to be $CO_2$, although the concept can be extended to other forcing agents. It is formally defined as the ultimate GMST warming that is realised after allowing the climate system to respond completely (the infinite time response) to an instantaneous doubling of $CO_2$. Within the forcing-response framework of equation 2.1.1, the ECS is defined as,

$$\mathrm{ECS} = \frac{F_{2\mathrm{X}}}{\lambda}, \tag{2.4.1}$$

where $F_{2\mathrm{X}}$ is the effective radiative forcing associated with a doubling of atmospheric $CO_2$ concentrations from their pre-industrial values (conventionally taken to be 280ppm). The ECS is clearly a function of both the radiative forcing of $CO_2$ and of the feedback of the system in response to the forcing. Uncertainty estimates in ECS combine uncertainties in both of these factors together.

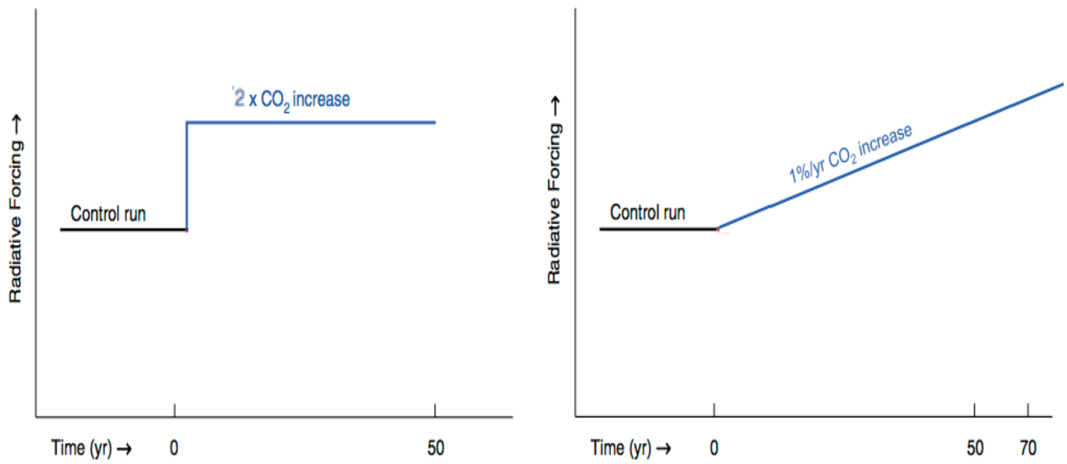


**Figure 2.5** – Schematic diagram for the idealised forcing profiles associated with the ECS (left) and TCR (right). Adapted from figures 1 and 2 of Kravitz *et al.* [2011].

The ECS is a fundamental metric of climate modelling and instantaneous $CO_2$ doubling experiments are useful for understanding the physics of atmospheric feedbacks and forcing. However, as the ECS represents the climate response on an infinitely long timescale, with concentrations held artificially fixed, it is debatable how directly relevant the ECS is to policy-analysis of anthropogenic climate change (although it is highly relevant to concentration stabilisation scenarios that remain the stated goal of the United Nations Framework Convention on Climate Change (UNFCCC) process), except in scenarios where ECS is very large and risks of runaway greenhouse effects may exist [Bloch-Johnson *et al.,* 2015]. A potentially more useful metric of climate forcing and feedback is the *transient climate response* (TCR). The TCR is defined as the GMST change averaged around the time of atmospheric $CO_2$ doubling as concentrations are increased by 1% each year (as the radiative forcing of $CO_2$ is logarithmic in the atmospheric $CO_2$ concentration, this corresponds to a linear ramp in forcing). The forcing profiles used to define ECS and TCR are shown in figure 2.5.

Within the forcing-response framework, the TCR can be expressed as,

$$\mathrm{TCR} = \frac{F_{2\mathrm{X}} - N_{2\mathrm{X}}}{\lambda}, \qquad (2.4.2)$$

where $N_{2\mathrm{X}}$ is the net energetic imbalance in the climate system at the time of atmospheric $CO_2$ doubling. It is clear from equation 2.4.2 that the energetic imbalance in the climate system masks the future warming potential of the forcing in the transient temperature response. The oceanic heat uptake can be thought of as a forcing on transient temperature change in a disequilibrium climate (e.g. Winton *et al.* [2010] and Held *et al.* [2010]), as could be seen via a rearrangement of equation 2.1.1 to leave $\Delta T$ on the left hand-side. Over the next few decades, and in scenarios where the radiative forcing continues to increase, or peaks and declines [Frame *et al.,* 2006; Allen *et al.,* 2009b], the long timescales and potential non-linear feedbacks [Andrews *et al.,* 2012] associated with ECS do not make ECS the most relevant measure of the climate response for the next century (again, noting the caveats about non-linearity and very large ECS given above). It has therefore often been argued that the TCR is more policy-relevant for contemporary climate change [Frame *et al.,* 2006; Hegerl, 2007].

Importantly, over recent years the incorporation of the carbon-cycle into Earth system modelling has allowed the impact of $CO_2$ emissions on the climate system to be more accurately quantified (and moved the modelling framework closer to something more directly applicable to quantifying the climatic impacts of policy decisions regarding $CO_2$ emissions). When including the carbon-cycle within models of the Earth system, the total cumulative emissions over all time emerges as the main determiner of $CO_2$-induced warming in the climate system [Allen *et al.,* 2009b; Meinshausen *et al.,* 2009]. This arises via a compensation between the logarithmicity of the relationship between radiative forcing and $CO_2$ concentrations (each additional increase in $CO_2$ concentration increases the radiative forcing on the climate system by less than the last) and the increase in cumulative airborne fraction with cumulative emissions (each additional unit of $CO_2$ emissions remains in the atmosphere longer than the previous one) driven primarily by non-linearities in the ocean carbonate chemistry (see e.g. Revelle & Suess [1957]). An approximately linear relationship between the cumulative emissions of $CO_2$ and the transient temperature response emerges from this compensation with a gradient defined as the *transient response to cumulative emissions* (TCRE - a measure of the peak warming expected for the emission of a trillion tonnes of carbon; see chapter 7), which in turn depends linearly on the TCR via,

$$\mathrm{TCRE} = \frac{\mathrm{TCR}}{\ln(2)} \ln\left(1 + \frac{\alpha_1 C_1}{C_0}\right), \tag{2.4.3}$$

where $\alpha_1$ is the cumulative airborne fraction (the fraction of a trillion tonne emission that remains in the atmosphere at the end of the injection) and $C_1$ and $C_0$ are the carbon loading of a trillion tonnes injection and the pre-industrial atmosphere respectively [Allen & Stocker, 2014]. Equation 2.4.3 would be expected to break down for very large ECS ($\lambda \to 0$ and equilibration time-scales $\to \infty$) where large fraction of the total $CO_2$-induced warming is only realised after concentrations have been stabilised.

## 2.5 Estimating climate forcing and response

ECS and TCR uncertainties can be investigated in GCMs and from observations. As the definitions of TCR and ECS (as shown diagrammatically in figure 2.5) involve idealised scenarios, GCMs can simulate these scenarios directly. Simulations that involve an instantaneous step in the $CO_2$ forcing (specifically a quadrupling) and a 1% annual increase in the $CO_2$ concentrations are core experiments of the international CMIP5 multi-model intercomparison project [Taylor *et al.,* 2012]. These simulations are conducted by different modelling centres around world giving a range of simulated TCR and ECS values [Andrews *et al.,* 2012]. In recent years, the IPCC has adopted the regression-based approach of Gregory *et al.* [2004] to derive the ECS and different components of the climate feedback parameter and effective radiative forcing [Myhre, 2013a; Flato, 2013]. This method computes a linear feedback parameter (and respective components) by finding the gradient between global radiative imbalance of the climate system and GMST warming as the model warms towards a new equilibrium state. The effective radiative forcing is therefore calculated as the intercept of the line when $\Delta T = 0$. The advantage of this method is that an estimate of the ECS can be made without having to run the model until equilibrium is reached, which can take millennia and be prohibitively expensive in GCMs and ESMs.

As an alternative to the GCM-based approach, by invoking a simple energy-balance model of the climate response to forcing it is possible to use observations of the energy budget over the historical period to probabilistically infer uncertainty intervals for TCR and ECS [Gregory *et al.,* 2002; Aldrin *et al.,* 2012; Lewis, 2013; Otto *et al.,* 2013]. Unlike idealised GCM simulations, methods involving observations are complicated by the multitude of forcing agents that are driving changes in the climate system over the historical period and limited through the need to assume a simple energy-balance representation of the climate response. However, a potential advantage of these methods is a more direct link to observations than in GCM-based methods. Observational estimates of ECS and TCR are explored in chapter 4.

## 2.6 Conclusions

The climate response to $CO_2$ radiative forcing is complex. Whilst major progress on the representation of the climate system in GCMs has been made over the last several decades, there remain substantial uncertainties over key components of the climate response and the underlying physical mechanisms of the relevant feedback loops. Throughout the rest of this thesis I will use the term 'climate response' as a generic term to refer to the GMST response to increased $CO_2$ concentrations, encapsulating both TCR and ECS. This is an alternative to using 'climate sensitivity' to capture both ECS and TCR, as this term is often used to refer exclusively to the ECS.

The greatest uncertainty in climate feedbacks can be ascribed to cloud-based feedback mechanisms, which are dependent on processes that are challenging to simulate in GCMs due to constraints on GCM grid spacing. Feedback mechanisms involving low clouds are especially relevant for determining the spread of the simulated ECS and TCR in climate models and can be especially difficult to constrain from observations. Whilst ultimately the most-likely route to improving understanding of these processes lies with increasing model resolution, the required resolution to better represent many important processes is still likely to be possible only far in the future. For the moment, perturbing the value of sub-gridscale parameters in GCMs and experimenting with different parametrisation schemes can offer insights into the diversity of possible feedback outcomes and physical mechanisms that could occur in the real world.

The next two chapters of this thesis build on the physics of the climate feedback mechanisms discussed in this chapter. They respectively address the two ways of investigating uncertainties in the climate response described in section 2.6. Chapter 3 creates a PPE of GCM simulations that aims to sample a very broad range of different climate responses through different realisations of the physics discussed in this chapter. Both idealised instantaneous $CO_2$ doubling and 1%/yr $CO_2$ increase experiments are conducted to understand more about ECS and TCR uncertainty within GCMs. Following this, chapter 4 estimates uncertainties in the climate response derived from observationally-based energy-balance methods before comparing these estimations of the climate response with the GCM estimates from chapter 3, before focusing further on often-neglected structural uncertainties in the estimation of TCR and ECS using energy-balance models constrained by observations.

# CHAPTER 3

# A perturbed physics ensemble of GCM simulations

Chapter 2 reviewed the physics underlying the climate response to radiative forcing. This chapter conducts idealised $CO_2$-forced experiments on a set of perturbed parameter versions of the HadCM3 GCM to examine uncertainties in features of the climate response. This perturbed physics ensemble (PPE) demonstrates a broad range of ECS and, novelly, TCR, underpinned by a range of different forcings and feedbacks across the ensemble. Section 3.1 discusses important background about the HadCM3 GCM and how parameters are perturbed within the GCM framework. Section 3.2 then describes sets of parameter combinations chosen in order to sample a wide range of possible climate responses whilst also maintaining radiatively-balanced control climates. Section 3.3 then goes on to describe several aspects of the climate response in this PPE to idealised $CO_2$-forced scenarios, that serve to motivate subsequent investigations in this thesis.

## 3.1 The HadCM3 GCM

The Met Office Hadley Centre Coupled Model Version 3 (HadCM3), is a GCM of the interacting atmosphere-ocean system [Gordon *et al.,* 2000]. HadCM3 is a version of the Met Office Unified Model for weather and climate prediction and runs with an atmospheric resolution of 2.5°x3.75° in the horizontal, with 19 levels on a hybrid coordinate in the vertical dimension. The oceanic model has a 1.25°x1.25° horizontal resolution with 20 levels of depth. The two components, along with sea-ice and land-surface models, are coupled without the use of prescribed flux adjustments that have historically been used to maintain radiative balance of the coupled climate system. The standard version of HadCM3 has a small radiative imbalance in its control state [Gordon *et al.,* 2000].

HadCM3, while dating back to the CMIP3 generation of GCMs, is still widely used and compares well with many of the current CMIP5 generation GCMs on many measures of simulation quality [Flato, 2013]. HadCM3 has proved to be a highly stable GCM under perturbed parameter experiments (in the sense that the GCM can be successfully integrated under large perturbations to its standard sub-gridscale parameters) [Sexton *et al.,* 2012; Yamazaki *et al.,* 2013], allowing very diverse simulated climates. These perturbed versions of the GCM can outperform the standard version of the GCM when examined in terms of specific variables [Williamson *et al.,* 2013].

To produce PPEs of HadCM3 I use the distributed computing resource of climate*prediction*.net, a citizen science project with around 40,000 volunteer computing nodes donated by members of the public from around the world [Allen, 1999]. A version of HadCM3 was produced for this project that can be run on a single processor. This involves a translation of the GCM code, from its original 64-bit version to a 32-bit version, so as to be usable by the diverse range of computers owned by the climate*prediction*.net participants. Additionally, during experiments for the later chapters of this thesis, I discovered a previously undocumented under-simulation of

Arctic sea-ice in the boreal summer present in the version of HadCM3 used for all perturbed physics experiments at the Met Office and with climate*prediction*.net. This bias arises from an inclusion of an interactive sulphur cycle in the GCM, which introduced the sea-ice bias when flux adjustments were removed in Vellinga & Wu [2008], although the bias wasn't noticed or documented at the time. Jackson *et al.* [2012] noticed the issue, but did not correct the model. The version of HadCM3 used in this thesis should correctly be considered a variant of the HadCM3 model, and would not be expected to simulate an identical climate to the CMIP3 HadCM3. A comparison of several fields between the two versions of the GCM shows small differences in control climate outside the high latitudes (not shown) and therefore, for simplicity, I refer to the standard parameters of the version of HadCM3 used in this thesis as 'standard HadCM3'.

### 3.1.1 Parametrisation schemes in HadCM3

HadCM3 has numerous parametrisation schemes for sub-gridscale processes. Some of the most important schemes are those representing convection and cloud-related processes [Rougier *et al.,* 2009]. This section aims to provide a brief summary of a couple of important parametrisation schemes in order to give a flavour of how parametrisation schemes, and the perturbed parameters within, influence the GCM. This section is not intended to provide a thorough and detailed description of the full workings of each parametrisation scheme (as there are many more than can be described in the space available here), which can be found in the underlying literature and model technical documentation, but rather acts to highlight examples of how important perturbations to sub-gridscale parameters affect the GCM physics.

**The convection scheme**

The convection scheme in HadCM3 is based on parcel theory, in which sub-gridscale convection and the resulting clouds are simulated as an ensemble of convective plumes which are averaged in order to provide the mean characteristics of the grid-box for the next dynamical timestep. The convection scheme is described in Gregory *et al.* [1996b]. The scheme works from the lowest atmospheric level upwards and tests on each layer whether a parcel lifted adiabatically to the next layer would be buoyant (defined as a potential temperature difference of greater than 0.2K) [Gregory & Rowntree, 1990]. If this condition is met, convection is initiated and the parcel continues to rise, mixing with the environmental air according to,

$$-\frac{\partial M_{pi}}{\partial p} = E_i - N_i - D_i, \tag{3.1.1}$$

where $M_{pi}$ is the mass flux of the convective parcel for convective plume $i$ in the ensemble, $E_i$ is the entrainment rate of environmental air into the rising parcel, $N_i$ is the rate of detrainment of parcel air into the environment, associated with turbulent mixing around the cloud edge and occurring at all levels of the assent, and $D_i$ is the forced detrainment that occurs only after the parcel has reached its level of neutral buoyancy. When the parcel reaches its level of neutral buoyancy, it is assumed that a sufficient proportion of the ensemble of convective plumes represented by the bulk parcel have detrained at that level such as to allow the parcel to just remain buoyant enough ($> 0.2$K potential temperature difference) in the next layer above to continue convective assent. As such, the convective parcel is used to represent an ensemble of convective plumes.

The entrainment and mixing rates ($E_i$, $N_i$ and $D_i$) are parametrised to be proportional to the convective parcel mass flux (e.g. $E_i = \epsilon M_{pi}$). The coefficient $\epsilon$ is made an explicit function of pressure to capture the variation of entrainment flux with cloud size (small clouds typically occur at high pressure and large, deep clouds at low pressure) according to,

$$\epsilon = k_{ent} A_E \frac{p}{p_*^2}, \tag{3.1.2}$$

where $A_E$ is a fixed constant, $p_*$ the reference surface pressure and $k_{ent}$ is the entrainment coefficient. It is the coefficient, $k_{ent}$, that is varied in parameter perturbations within the climate*prediction*.net ensemble as the *entcoeff* parameter, and is an important determiner of cloud and radiative properties in the climate system.

### Cloud scheme parametrisations

Prognosed water content variables are required in GCMs to ensure that all liquid and frozen water generated by the resolved-scale motions is not precipitated out instantaneously when formed. HadCM3 treats frozen and liquid water precipitation processes separately [Gregory *et al.,* 1996a]. In liquid water clouds, the evolution of the liquid water content, $q_{lw}$, when condensing onto cloud-condensation nuclei, is given as,

$$\frac{\partial q_{lw}}{\partial t} = -C\left(c_T\left(1 - \exp\left[-\left(\frac{\rho q_{lw}/C}{c_w}\right)^2\right]\right) + \frac{P}{c_a}\right)\frac{q_{lw}}{C}, \tag{3.1.3}$$

where $C$ is the fraction of the grid-box covered by clouds, $c_T$ is a constant term representing the rate at which cloud water is converted to rain and the term $\frac{P}{c_a}$ is the accumulation of water from the cloud onto falling precipitation ($P$) droplets. Liquid precipitation is computed as the loss of the cloud liquid water content in order to conserve moisture. The parameter $c_T$, the rate of conversion of cloud water to precipitation, is varied in the climate*prediction*.net ensemble as the parameter *ct.* The threshold for precipitable water $c_w$ (the rate constant for the formation of precipitation from condensation onto cloud condensation nuclei), is also perturbed in the climate*prediction*.net ensemble as the parameter *cw-land* and *cw-sea* over land and ocean respectively (concentrations of cloud condensation nuclei differ between land and ocean).

In the frozen water precipitation scheme the rate of loss of frozen water content from the cloud is dependent on the speed of falling frozen precipitation, $v_f$. In HadCM3 this is determined via the parametrisation,

$$v_f = v_{f1} \left[ \frac{\rho q_c / C}{c_f} \right]^{0.17}, \tag{3.1.4}$$

where $v_{f1}$ is a scaling factor and the terms within the square brackets encode the dependencies on the environment. In the climate*prediction*.net ensemble this scaling factor is varied as the parameter *vf1*.

The parametrisation schemes and coefficients described above are just a couple of the many parameters that are perturbed in climate*prediction*.net ensembles of HadCM3 (see table 3.1). For details of how other perturbed parameters are related to the GCM physics, interested readers should consult the Unified Model documentation.

## 3.2 A radiatively-balanced non-flux-adjusted perturbed physics ensemble

An advantage of the climate*prediction*.net distributed computing infrastructure is the capability to create very large PPEs. Previous ensembles run under climate*prediction*.net have contained on the order of tens of thousands of ensemble members (e.g. Rowlands *et al.* [2012]). In this chapter (and throughout this thesis) 32 independent parameters, from both the atmosphere and ocean models, are perturbed. Table 3.1 lists these parameters and their reference values.

Whilst flux adjustments in non-perturbed GCMs have not been common since the early 2000s, perturbations to the parameters within GCMs has previously disrupted the radiative balance of the GCM control states and so required the re-introduction of radiative flux adjustments in order to preserve radiative balance in PPEs [Sanderson *et al.*, 2008]. Yamazaki *et al.* [2013] attempted to solve this problem by using a very large ensemble of HadCM3 perturbed parameter experiments that were run without flux adjustments to predict a region of parameter space consistent with a control climate in net radiative balance (within the uncertainties of the standard HadCM3 pre-industrial climate state).

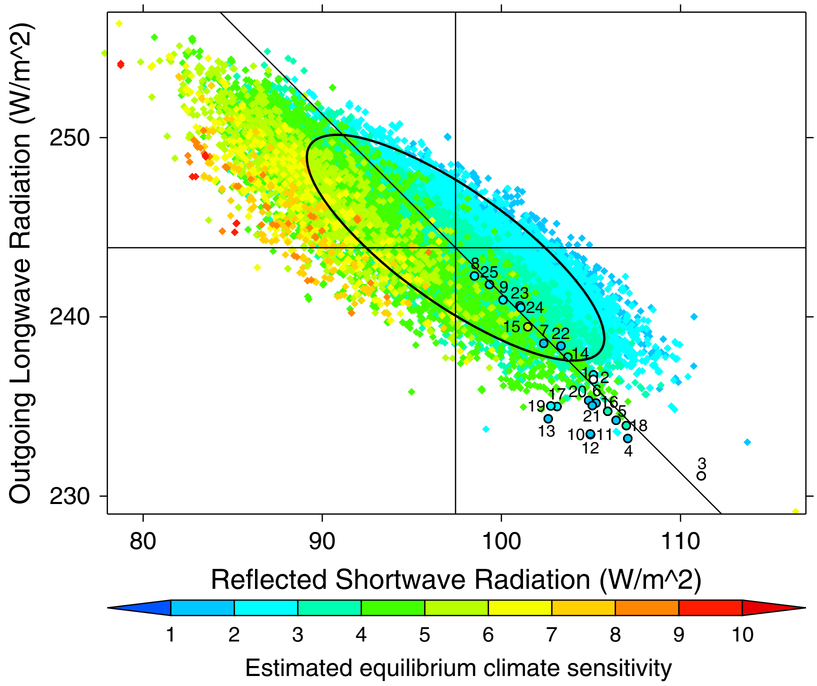


**Figure 3.1** – Figure 4 from Yamazaki *et al.* [2013]. A 'flux-balanced' ensemble of parameter perturbations are shown, coloured by their predicted ECS based on an emulation using only their atmospheric parameter perturbations (with the emulator having been trained using a previous climate*prediction*.net slab GCM ensemble). The black ellipse represents the 99% confidence interval in the total observational-model discrepancy of TOA flux observations (see Tett *et al.* [2013] for details) centred around the standard physics HadCM3 configuration (intersection of the straight black lines). The open black circles show the other CMIP3 GCMs. The diagonal indicates a slope -1:1 line through the standard HadCM3.

Figure 3.1, a reproduction of figure 4 from Yamazaki *et al.* [2013], shows PPE members that are targeted to have TOA radiative fluxes within the 99% confidence interval of total TOA observational-model uncertainty from Tett *et al.* [2013] (black ellipse), centred around the standard physics HadCM3 configuration. Whilst not all members of the ensemble are found to lie exactly within the black ellipse, some GCMs are emulated to have TOA fluxes that are very close to the standard

HadCM3. The range of reflected short-wave radiation (RSR) and outgoing long-wave radiation (OLR) compares favourably to the CMIP3 MME, shown as open black dots in figure 3.1.

Yamazaki *et al.* [2013] described the pre-industrial control climate of these GCMs. In this chapter I use these same parameter sets as a first ensemble to examine uncertainty in the climate response under idealised $CO_2$ forcing scenarios. The analysis in the next section documents the range of climate response in this PPE as well as serving as an introduction to the GCM analysis methods used across the rest of the thesis.

## 3.3 Climate response in a balanced perturbed physics ensemble

In this section I document the climate response of the Yamazaki *et al.* [2013] parameter sets under both a constant $CO_2$ forcing as well as a steadily increasing $CO_2$ forcing. In what I subsequently refer to as the STEP experiment, I perturb the climate with an instantaneous doubling of $CO_2$ concentrations after 10 years of pre-industrial conditions with concentrations held constant thereafter. In what I call the RAMP experiment I annually increase atmospheric $CO_2$ concentrations by 1%, with concentrations reaching double pre-industrial levels after 70 years. Throughout this chapter the range of response from the PPE is contrasted with the response of an initial condition ensemble of the standard HadCM3, in which numerical noise is added to the potential temperature field of the atmospheric start conditions, under the same STEP and RAMP forcings. These ensembles will be known as the STEP-init and RAMP-init ensembles respectively.

Due to the bandwidth and data storage limitations of the climate*prediction*.net project, and the size of the ensembles, spatial fields are recovered only for the seasonal averages in each decade of each ensemble member. Whilst STEP and RAMP

**Table 3.1** – HadCM3 perturbed parameters in climate*prediction*.net (based on tables in Yamazaki *et al.* [2013]).

| Name | Description | Default value |
|---|---|---|
| CT | Clouds- Conversion rate of liquid water to precipitation | $1 \times 10^{-4}$ ($s^{-1}$) |
| CW_SEA | Clouds- Threshold cloud liquid water content over sea | $5 \times 10^{-5}$ ($kgm^{-3}$) |
| CW_LAND | Clouds - Threshold cloud liquid water content over land | $2 \times 10^{-4}$ ($kgm^{-3}$) |
| EACF | Clouds - Cloud fraction at grid-box saturation | 0.5 in 19 levels |
| RHCRIT | Clouds - Critical relative humidity for cloud fraction | 0.95,0.9,0.85,0.7 in 16 levels |
| VF1 | Clouds - Ice fall speed | 1 ($ms^{-1}$) |
| ENTCOEF | Convection - Entrainment rate coefficient | 3 |
| ALPHAM | Radiation - Albedo at melting point of sea ice | 0.5 |
| DTICE | Radiation - Temperature over which ice albedo varies | 10K |
| ICE_SIZE | Radiation - Ice particle size | $30 \times 10^{-6}$ (m) |
| DIFF_COEFF | Dynamics - Horizontal diffusion coefficient for temperature and wind speed | $5.47 \times 10^{8}$ in 18 levels, $4 \times 10^{6}$ |
| DIFF_COEFF_Q | Dynamics - Horizontal diffusion coefficient for water vapour | $5.47 \times 10^{8}$ in 13 levels, $1.5 \times 10^{8}$ in 5 levels, $4 \times 10^{6}$ |
| DIFF_EXP | Dynamics - Order of horizontal diffusion for temperature and wind speed | 3 in 18 levels, 1 |
| DIFF_EXP_Q | Dynamics - Order of horizontal diffusion for water vapour | 3 in 13 levels, 2 in 5 levels, 1 |
| KAY_GWAVE | Dynamics - Surface gravity wave drag: typical wavelength | $2 \times 10^{4}$ (m) |
| KAY_LEE_GWAVE | Dynamics - Surface gravity wave trapped lee wave constant | $3 \times 10^{5}$ ($m^{-3}$) |
| START_LEVEL_GWDRAG | Dynamics - Lowest model level for gravity wave drag | 3 |
| R_LAYERS | Land Surface - Number of soil levels from which water can be extracted | 4,4,3,3,7 |
| ASYM_LAMBDA | Boundary Layer - Vertical distance over which air parcels travel before mixing | 0.15 |
| CHARNOCK | Boundary Layer - Constant for calculating roughness length for momentum transport over sea | $1 \times 10^{-2}$ |
| G0 | Boundary Layer - Calculation of stability function for heat, moisture, and momentum transport | 10 |
| Z0FSEA | Boundary Layer | $1 \times 10^{-3}$ (m) |
| ISOPYC | Ocean - Isopycnal diffusion of tracer at surface | 1000 ($m^2s^{-1}$) |
| VDIFFSURF | Ocean - Background vertical diffusion of tracer at surface | $1 \times 10^{-5}$ ($m^2s^{-1}$) |
| VDIFFDEPTH | Ocean - Increase of background diffusion of tracer with depth | $3 \times 10^{-8}$ ($ms^{-1}$) |
| VERTVISC | Ocean - Background vertical diffusion of momentum (viscosity) | $1 \times 10^{-5}$ ($m^2s^{-1}$) |
| MLDEL | Ocean - Decay of wind mixing energy with depth | 100 (m) |
| MLLAM | Ocean - Wind mixing energy scaling factor | 0.1 |

experiments were commenced for all of the Yamazaki *et al.* [2013] parameter sets, only a subset of these were returned to the climate*prediction*.net servers. The failure of a simulation can be caused by many potential issues on a participant's computer, such as an abrupt shut-down or crash of the machine. Other computing issues can also cause corrupted or missing data and therefore I parse the raw ensemble data to identify such issues and remove those simulations before conducting the analysis of the ensembles presented in this thesis. After additionally removing ensemble members that seem to be clearly running away to infinite warming (based on an inspection of the temperature and TOA radiation fields) 1241 independent parameter sets were included in analysis of the STEP experiment and 761 were included in analysis of the RAMP experiment. The subset of parameter sets that are common to both is 711. The initial condition ensemble experiments were completed by 71 members. In this section, all quantities are expressed as differences from the the timeseries of the corresponding control experiment.

The subsections below focus in turn on different aspects of the climate response across the ensemble. Namely, components of forcing and feedback and the relationship between the STEP and RAMP experiment responses including the variation in spatial patterns of change across the ensemble.

### 3.3.1 Forcings and feedbacks

Figure 3.2 shows the GMST and global-mean precipitation response under the STEP and RAMP experiments. The diversity of climate response across the ensembles is clear. The ensemble mean warms by approximately 2.34K after 70 years of doubled $CO_2$ concentration (almost identical to the standard HadCM3) and the 5-95 percentile spread of the ensemble response is 1.79-3.34K. Under both STEP and RAMP experiments, precipitation increases more in the ensemble mean than the standard HadCM3 despite GMST temperatures being very similar. The density of ensemble members is largely concentrated at the mid-low end of the response

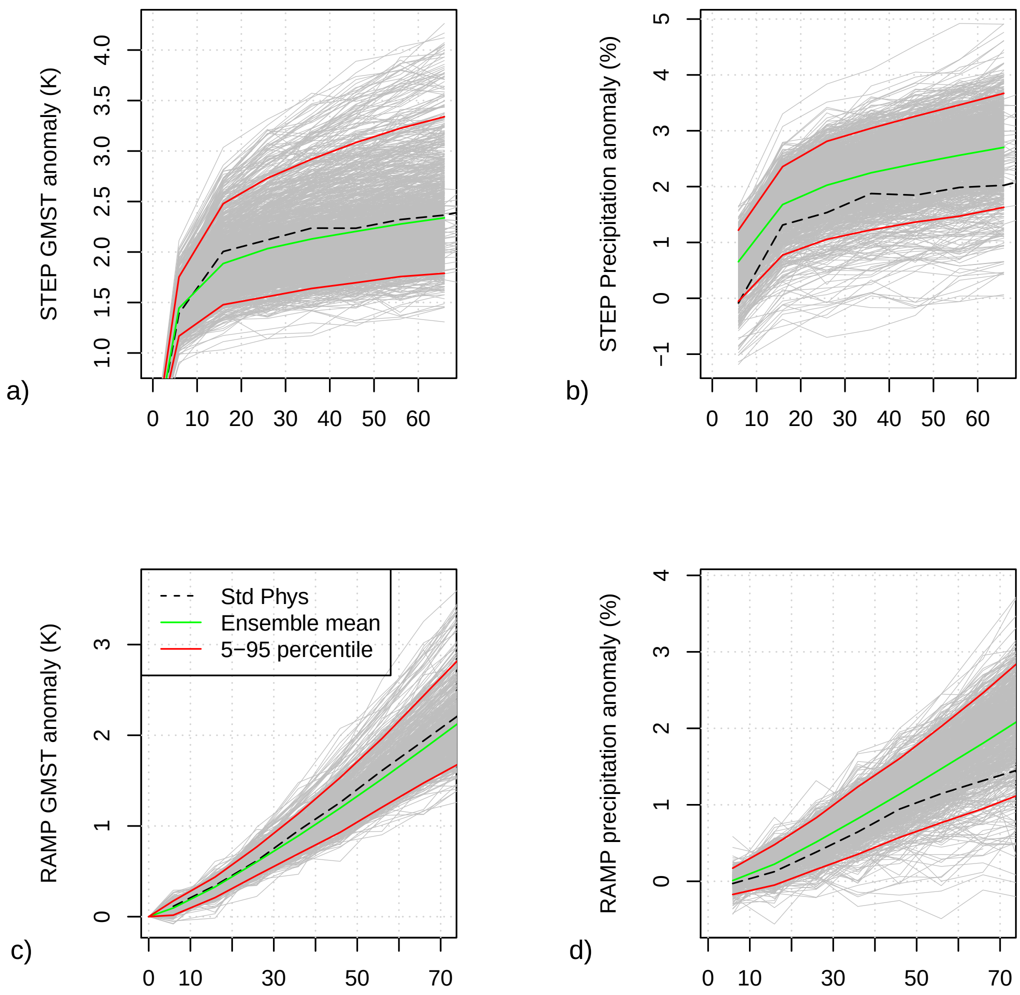


**Figure 3.2** – The evolution of GMST (panels a) and c)) for the STEP and RAMP ensembles respectively. Percentage change in the global-mean precipitation is shown for the STEP and RAMP ensembles in panels b) and d) respectively. The x-axis shows the years since doubling for the STEP experiment and since the start of the forcing increase in the RAMP experiment. Grey lines show ensemble members and the broken black line shows the ensemble mean of the initial conditions standard physics ensemble. The green line denotes the STEP/RAMP ensemble means and the red lines the 5-95 percentiles.

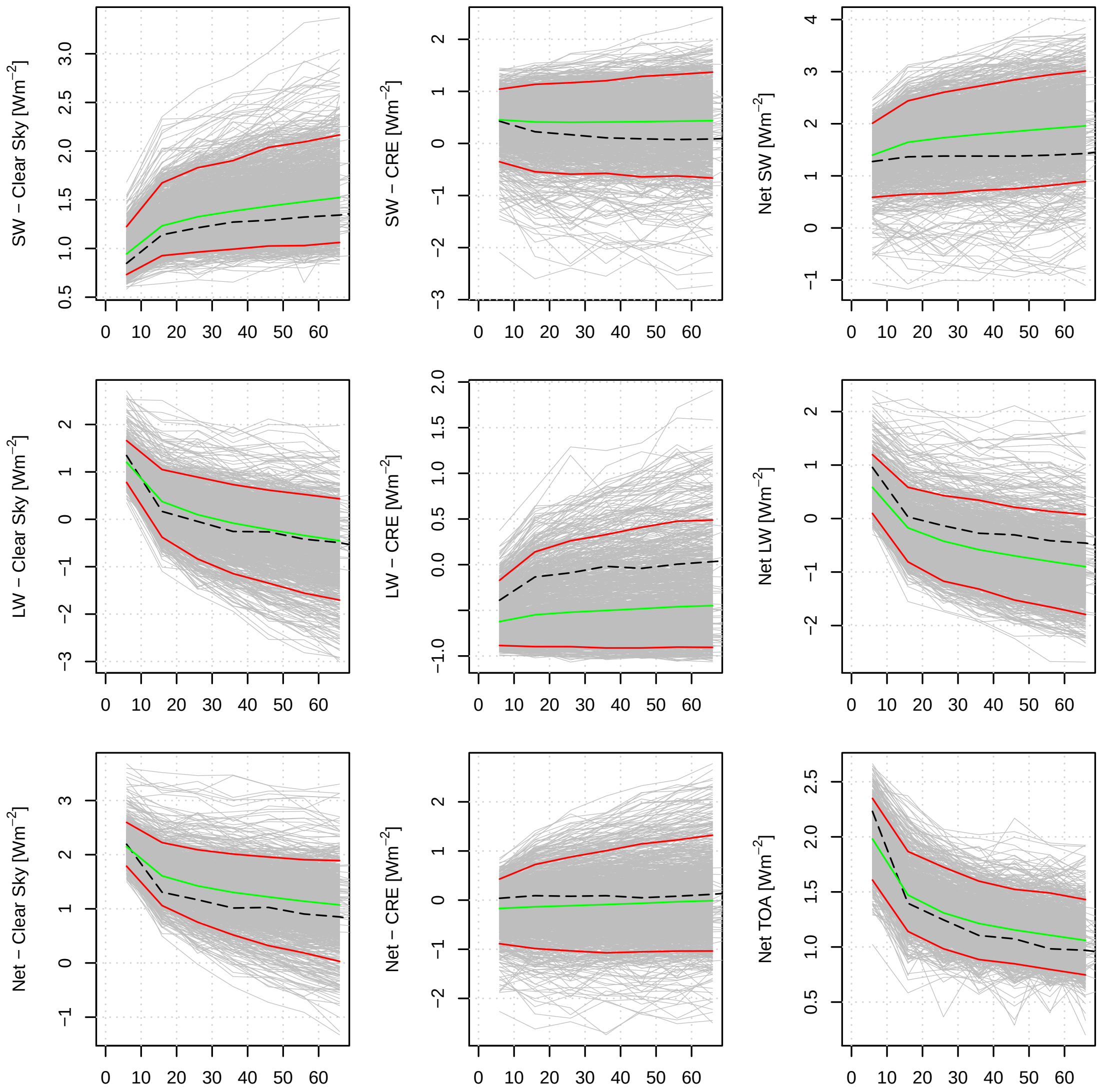


**Figure 3.3** – As for figure 3.2, but for the evolution of the components of the TOA radiative flux in the STEP ensemble.

range with a long-tail of high response ensemble members.

Similarly, a large range of responses in the planetary energy imbalance under the STEP experiment is clear from figure 3.3. In figure 3.3 all fluxes are defined as positive downwards and the cloud radiative effect (CRE) terms are defined as the difference in radiative fields between conditions with clouds present and with clouds artificially removed. Ensemble members exist in which short-wave (SW) and long-wave (LW) CRE increase and decrease over time with ensemble members also displaying both signs of the instantaneous fast adjustment of the CRE on imposition of the radiative forcing.

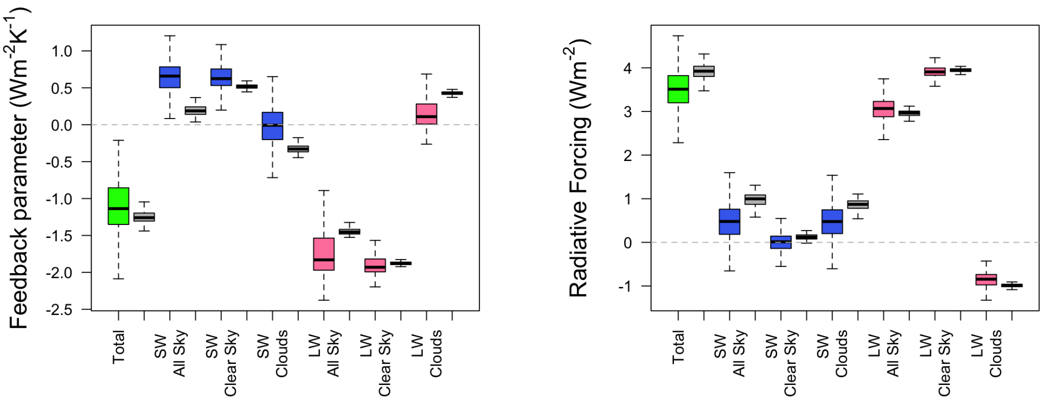


**Figure 3.4** – The range of feedback components (left) and components of effective radiative forcing (right) as simulated in the STEP ensemble. Forcing and feedback components are derived via regressions of components of TOA energetic imbalance against the GMST anomaly. For each component, the range from the initial condition ensemble are shown as grey box-plot adjacent to the respective component.

The now-typical approach to deriving climate feedbacks computes the components of forcing and feedback from the intercept and slope of a regression of a component of TOA imbalance against GMST change under a fixed forcing [Gregory *et al.,* 2004]. Figure 3.4 shows the components of global climate feedback as calculated through linear regressions of the TOA flux components against changes in GMST. The PPE (coloured bars) simulates a very broad range in all components of the climate feedback parameter. Parameter perturbations within HadCM3 most strongly affect cloud components of the radiative feedbacks, they also have an effect on the magnitude of the clear-sky feedbacks, albeit much smaller, associated with differing sea-ice, snow cover and atmospheric lapse-rate change under warming. In all components of climate feedback the PPE spans ranges several times larger than the ranges associated with internal variability and the uncertainty in the regression-based method (grey bars). Particularly in CRE feedbacks, the PPE displays considerably different behaviour to its standard-configuration counterpart with both LW and SW CRE feedback ranges spanning both positive and negative values. The ensemble median value for SW cloud feedback sits very close to the zero feedback line and the range of SW-CRE feedbacks is relatively evenly distributed either side of zero. On average the ensemble has a more positive SW-CRE feedback and a more negative LW-CRE feedback than the standard HadCM3, indicating that the perturbed ensemble members in general have clouds that are less responsive to warming. Contrasting the range of feedbacks in the PPE to the range of feedbacks in the CMIP5 MME (net feedback parameter ranges are calculated to be -2.3 to -0.9 $\mathrm{Wm}^{-2}$) shows a larger spread of the total feedback component in the PPE, again largely driven by the spread in the simulated cloud feedback components. A comparison of the median values of the feedback components in figure 3.4 to the median values of the standard physics ensemble offers a useful reminder that the PPE ensemble mean (or median) should not be interpreted as being equivalent to the response of the the standard HadCM3.

Sanderson *et al.* [2008] examined the contribution of different parameter perturbations to the overall feedback parameter in a previous ensemble of the HadSM3 GCM (which has a similar atmospheric model to HadCM3 but is coupled to a slab or thermodynamic mixed-layer ocean instead of the full-depth dynamic ocean used here) in which single atmospheric parameters were perturbed discretely. Whilst that GCM and experimental structure is not identical to the one here, within a simple linear regression framework similar parameters are found to be the most influential in determining the components of climate feedback, namely *entcoef* (the entrainment coefficient), *vf1* (the ice fall speed) and the *ice_size* (the radius of ice particles as seen by the radiation scheme) parameters explain the greatest variance in net climate feedback parameter (not shown). However, in an ensemble with 32 independently and continuously perturbed parameters the dependences of the components of climate feedback are necessarily complex and highly multi-dimensional, with non-linear terms and interaction effects between parameters expected to play an important role. A full statistical emulation of the effect of specific parameter perturbations on aspects of the climate response of HadCM3 is left until chapter 5. The STEP ensemble described here also varied oceanic parameters within the HadCM3 dynamical ocean model. Whilst oceanic parameter perturbations have previously not been found to as important to the climate response as atmospheric perturbations [Collins *et al.*, 2007], here I find that variations in the coefficient of background vertical diffusivity in the upper ocean, *vdiffsurf*, is the fifth most important variable (not shown) in explaining variations in the global climate feedback parameter and is more important for global climate feedbacks than the mixed layer parametrisation scheme perturbations, as found in Collins *et al.* [2007].

Estimates of the ECS depend on estimates of the effective radiative forcing of $CO_2$ as well as the climate feedback parameter. Figure 3.4 (right hand panel) shows the ranges of the different components of radiative forcing across the STEP ensemble. Large ranges of different effective radiative forcings are simulated for SW-CRE, indicating variance within the ensemble of rapid adjustments to the overall cloud reflectivity as a direct consequences of the enhanced $CO_2$ in the atmosphere. In contrast, the range of rapid-adjustments to the LW-CRE is much more constrained. Clear-sky effective radiative forcings are more similar to the standard HadCM3 than their corresponding feedbacks, offering encouragement that the diversity of ECS in the STEP ensemble is not generated by drastically altering physically well-understood clear-sky forcing of $CO_2$. SW-CRE fast adjustments are the main determiner of the spread of effective radiative forcing within the ensemble, with SW-CRE fast adjustments having both heating and cooling effects in different ensemble members. LW-CRE and SW-CRE fast adjustments are weakly ($r^2 = 0.11$) anti-correlated across the ensemble, indicating a certain level of coupling between the underlying physical processes driving the adjustments.

The shape of the histogram of ECS in figure 3.5 shows a characteristically skewed shape with a long-tail of a small number of high sensitivity ensemble members. The range of ECS (1.80-9.63K) across the ensemble is driven more by variations in the net climate feedback parameter as opposed to the effective radiative forcing of $CO_2$. In contrast, the shape of the TCR distribution in the right hand panel of figure 3.5 shows a more symmetrical shape, albeit still with some positive skew, similar to other TCR distributions typically used in the literature [Rogelj *et al.,* 2012; Harris *et al.,* 2013]. Fixing effective radiative forcing at the ensemble average value, $F_{ave}$ (3.50Wm$^{-2}$), gives an ECS range (calculated as $\frac{F_{ave}}{\lambda}$) of 1.6 - 15.9K due

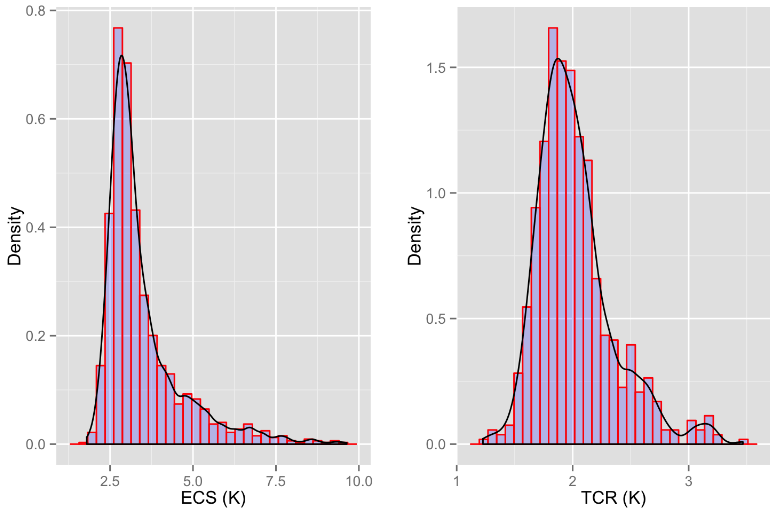


**Figure 3.5** – Histograms of ECS (left) and TCR (right). ECS is estimated from the forcing and feedback components in the STEP ensemble (derived through the linear regression method). The black line shows a probability density estimate from a kernel density estimator (a non parametric method of estimating a probability density function).

to variations in the feedback parameter alone, greater than the ECS range for the net feedback parameter fixed at the ensemble mean value, $\lambda_{ave}$ (-1.10Wm$^{-2}$K$^{-1}$), and varying effective radiative forcing (1.6-4.8K). As the range of ECS for fixed forcing is greater than the actual range of ECS across the ensemble, this indicates evidence of anti-correlation between effective radiative forcing and feedback parameters, as observed in the CMIP5 MME [Andrews *et al.*, 2012; Ringer *et al.*, 2014]. Across this ensemble effective radiative forcing and total feedback parameter are anti-correlated with correlation coefficient -0.73. The number of simultaneously changing variables between the ensemble members makes it difficult to tease out a mechanistic understanding of this inter-dependence of forcing and feedback, however, the correlation is observed to be the strongest in the clear-sky relationships and weakest in the cloud radiation fields (particularly SW-CRE, where $r = -0.43$), suggesting that cloud rapid-adjustments and warming-driven feedbacks occur more as decoupled independent processes, whilst a fundamental physical connection may underlie the coupling between clear-sky forcing and feedback.

Slab ocean configurations of GCMs have traditionally been used for the computation of ECS, as they could inexpensively be run to equilibrium, unlike full 3D

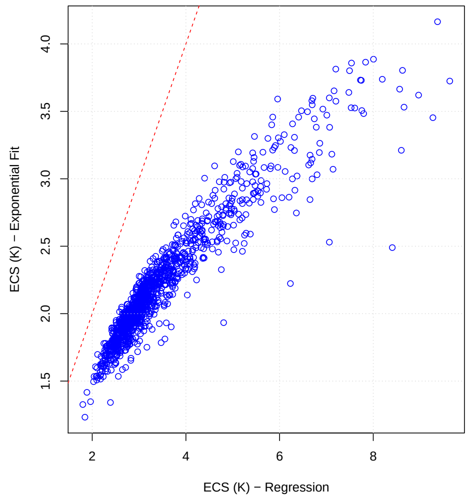


**Figure 3.6** – A comparison of the ECS estimated through the regression-based method (x-axis) and through a fit of the GMST anomaly timeseries with a single exponential model (y-axis). The red dashed line shows 1:1.

oceans in which ocean circulation takes many thousands of years to adjust to the forcing, delaying the full equilibration of surface temperature accordingly [Held *et al.*, 2010]. Additionally, it has long been known that lower ECS GCMs approach equilibrium faster than high ECS GCMs [Manabe & Wetherald, 1967]. Slab ocean ECS has often been derived via fitting exponential adjustment statistical models to the temperature anomaly timeseries (e.g. Stainforth *et al.* [2005]). Figure 3.6 contrasts the estimation of the ECS from the regression method with an estimate of the ECS derived from a fit of the GMST anomaly timeseries with a single exponential adjustment model of the form, $T(t) = ECS(1-\exp[-kt])$ for the full 3D ocean GCMs considered here. When calculated via an exponential fit, the ECS is consistently underestimated when compared to the the regression intercept method. At the high sensitivity end, many of the GCMs in the STEP ensemble still have substantial (0.5Wm$^{-2}$ or greater), albeit slowly evolving, net TOA flux imbalances, a likely reason for the considerable underestimate of the ECS through the exponential fit. At lower sensitivities, the exponential fit and the regression methods are much more similar, but still have differences of the order of 0.5K between the methods. The regression-based method is known [Gregory & Webb, 2008] to perform worst for GCMs at the low sensitivity end of the range, as natural variability fluctuations can be a substantial fraction of the $CO_2$-induced signal, however figure 3.6 shows that even in low response cases substantial energetic imbalances can still exist after 70 years of doubled concentrations and deriving ECS instead through exponential fits is still likely to produce a biased low answer for all 3D ocean GCMs.

## 3.4 Relationship of the RAMP response to the STEP response

In the STEP experiment radiative forcing is held fixed and the temporal evolution of the climate solely represents the changing disequilibrium state of the climate system. TCR in the RAMP ensemble spans the range 1.23-3.47K. However, in the RAMP experiment there is an additional complexity of a time-varying radiative forcing. For the standard version of HadCM3 [Good *et al.*, 2011], as well as other GCMs in the CMIP5 MME [Good *et al.*, 2013], it has been demonstrated that the response of the GCM to linearly increasing forcing can be successfully estimated from a linear step-response scaling of the response under the instantaneously doubled forcing. This method allows the response to an arbitrary $CO_2$ forcing scenario to be computed from a scaled summation of the response to a much simpler forcing scenario. Whilst this has been shown to be a good approximation for some variables in the standard HadCM3, it remains unclear how generally valid this linear step-response approximation is across a range of different climate sensitivities and GCM parameters.

The linear step-response model from Good *et al.* [2011] is mathematically described as,

$$y_i = \sum_{j=0}^{i} x_j \frac{\Delta f_{i-j}}{\Delta f_s}, \tag{3.4.1}$$

where $y_i$ is a geophysical variable of interest in year $i$ of a transient forcing simulation, $x_j$ is the same variable in year $j$ of the step forcing simulation and $\frac{\Delta f_{i-j}}{\Delta f_s}$ is the ratio of the change in radiative forcing in year $i-j$ of the transient forcing simulation to the change in radiative forcing in the step experiment. Moving from the general case to the specific case of the RAMP ensemble, equation 3.4.1 can be expressed as,

$$y_i = \sum_{j=0}^{i} x_j \frac{1}{7}, \tag{3.4.2}$$

as the RAMP scenario is represents a linear increase in forcing and the temporal frequency is now the decadal mean output from the STEP and RAMP ensembles. Physically, the step-response model decomposes the climate response in each decade under the RAMP forcing into the summation of the climate response under a 1/7th of the STEP experiment instantaneous increase in $CO_2$ concentrations in each of the preceding decades.

Figure 3.7 shows the ability of the step-response model to accurately predict the response of a member of the RAMP ensemble (averaged over years 60-80 of the forcing ramp) from the corresponding member of the STEP ensemble for a number of different variables of interest. For all the variables shown, and for all others examined, the step-response emulation faithfully captures the evolution of the simulated RAMP experiment across the large parameter perturbations within the ensemble at the global and decadal mean level, indicating that a linear relationship between the climate response to fixed and changing radiative forcing is robust across a wide range of ECS and GCM parameter perturbations. The step-response model consistently marginally underestimates the GMST response in the RAMP ensemble, and underestimates precipitation changes for higher ECS GCMs, indicating the possibility of non-linear physics in driving high ECS, similar to as highlighted in Good *et al.* [2013]. Deviations from linearity across the ensemble are not statistically significantly (at the 5% level) associated with any particular parameter perturbations within the ensemble. Planetary heat uptake is particularly well predicted via the linear step-response model, indicating a limited impact of evolving patterns of oceanic heat uptake in the STEP integrations.

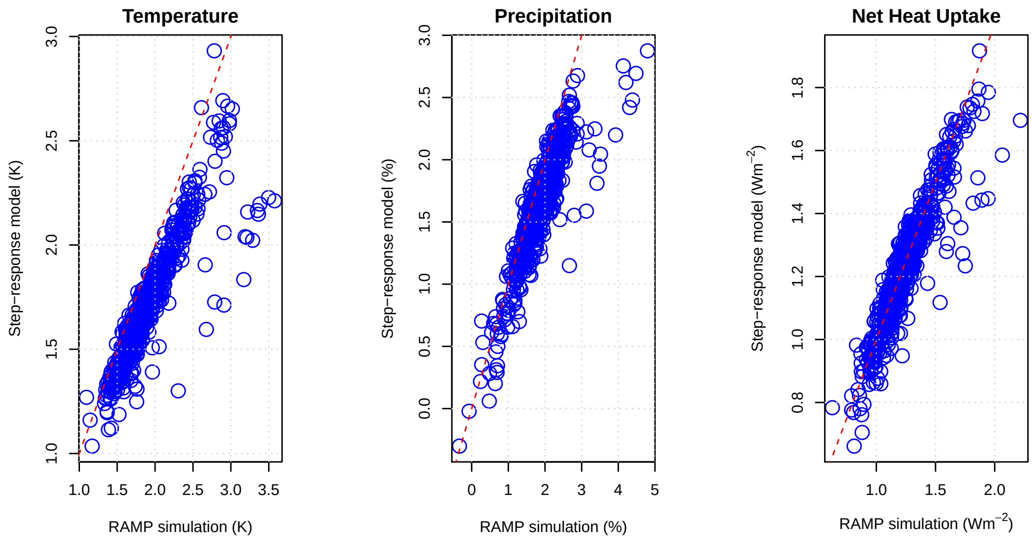


**Figure 3.7** – Step-response emulation of the RAMP ensemble from the STEP ensemble (y-axis) compared to the simulated value in the RAMP ensemble (x-axis), the red line shows the 1:1 relationship. Relationships are shown for GMST change (left), precipitation change (middle) and net planetary heat uptake (right) respectively.

### 3.4.1 Variations in the pattern of warming across the perturbed physics ensemble

Several previous studies have tried to quantify the effects of perturbing GCM parameters on the magnitude of warming, particularly in the global-mean [Murphy *et al.*, 2004; Stainforth *et al.*, 2005; Rougier *et al.*, 2009; Sexton *et al.*, 2012]. However, looking beyond the global-mean, there has been relatively little previous attention to the impact of parameter perturbations on the *pattern* of warming for a given increase in GMST. This section briefly summarises the impact of parameter perturbations on patterns of surface temperature change across the STEP ensemble and differences between the surface warming pattern associated with a given GMST increase between the STEP and RAMP ensembles.

To define the pattern of warming for a unit increase in GMST, the fields of surface temperature change in the STEP and RAMP ensembles are decomposed into a scaling factor associated with GMST change and a spatial pattern [Mitchell *et al.*, 1999; Huntingford & Cox, 2000]. A generalised pattern-scaling, $\chi(x, y)$, of a climatic variable is calculated as,

$$\chi(x,y) = \frac{\int \bar{S}(t)X'(t,x,y)dt}{\int \bar{S}(t)^2 dt}, \tag{3.4.3}$$

where $X'(t,x,y)$ is the spatially and temporally varying geophysical field of anomalies from the control simulation and $\bar{S}(t)$ is the spatially averaged global anomaly of the scaling factor. The best-estimate GCM response when the scaling factor anomaly has value $\bar{S}^* = \bar{S}(t^*)$ is given by,

$$X(t^*,x) = \bar{S}^*\chi(x,y). \tag{3.4.4}$$

To derive the best-estimate pattern-scaling for each ensemble member, I use GMST anomaly as the scaling factor with the pattern $\chi(x,y)$ derived via a best-estimate gradient of the linear regression between $X'(t,x,y)$ and $\bar{S}(t)$ for data points representing each decade of the STEP and RAMP ensemble members respectively.

Modes of variance that account for significant fractions of the variation in these derived pattern-scalings across the ensemble can be found through an empirical orthogonal function decomposition (EOF). This technique identifies patterns corresponding to orthogonal axes of variation across the ensemble (EOFs) as well as principal components, that indicate how much of each EOF is expressed within each ensemble member. Following the method of Sanderson *et al.* [2008] it is possible to relate the variations in the principal components of a specific EOF across the ensemble with the variations in the underlying parameter perturbations. Constructing a linear model, the principal components (PCs) can be related to the underlying parameter perturbations through,

$$P_{ij} = \sum_{i=1}^{n} p_{il}\beta_{lj} + noise, \tag{3.4.5}$$

where $P_{ij}$ are the normalised (anomaly from the ensemble mean divided by the ensemble standard deviation) principal components of the $j^{th}$ EOF in the $i^{th}$ ensemble member, $p_{il}$ is the normalised value of parameter $l$ in the $i^{th}$ member of the ensemble and $\beta_{lj}$ is a regression coefficient expressing how strongly variations of parameter $l$ across the ensemble influences the expression of the $j$th EOF in the ensemble members.

Figures 3.8 and 3.9, show the first three EOFs of pattern-scaling variation across the PPE in the RAMP and STEP experiments respectively. The first two dominant modes of variation are very similar between the RAMP and STEP experiment and explain almost identical fractions of the total variance across the ensemble. The dominant mode of variability in both ensembles is the relative strength of tropical warming relative to high latitude warming for a given global-mean warming, likely to be associated with the differing cloud and convective feedbacks across the ensemble which would be expected to have most impact on the tropics. A particularly pronounced signal is seen over the Amazon basin, possibly associated with differing land-surface responses in this region. The second mode of variance contains strong signals in regions of deep water formation in the North Atlantic and upwelling (particularly in the Southern Ocean) and is likely to be associated with the diverging responses of the global overturning circulation that exist within the ensembles. Whilst similar dominant modes of pattern-scaling variability are apparent within the RAMP and STEP ensembles, the linear regression of the principal components onto GCM parameters as calculated via equation 3.4.5, associates these modes of variability with very different parameter perturbations in the STEP and RAMP ensembles (not shown), none of which have statistically significant regression coefficients, indicating that parametric drivers of this variability cannot be deduced from the experimental set-up used here under a simple linear regression analysis.

## 3.5 Summary

This chapter had two aims; firstly to introduce the experimental methods of PPEs that are employed in subsequent chapters of this thesis and secondly to describe

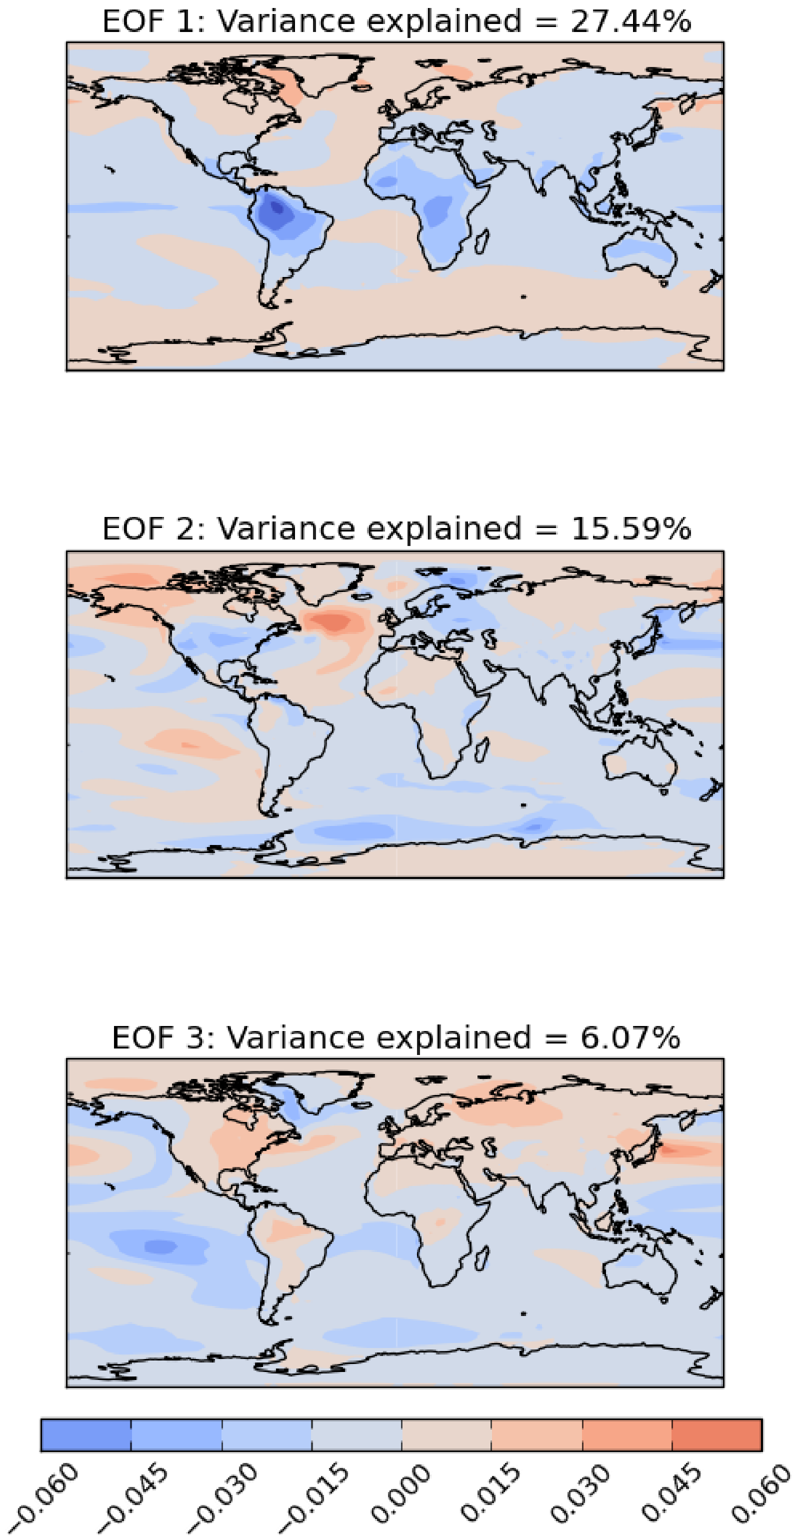


**Figure 3.8** – EOFs of the dominant three modes of variation in the pattern-scaling of warming across the STEP ensemble.

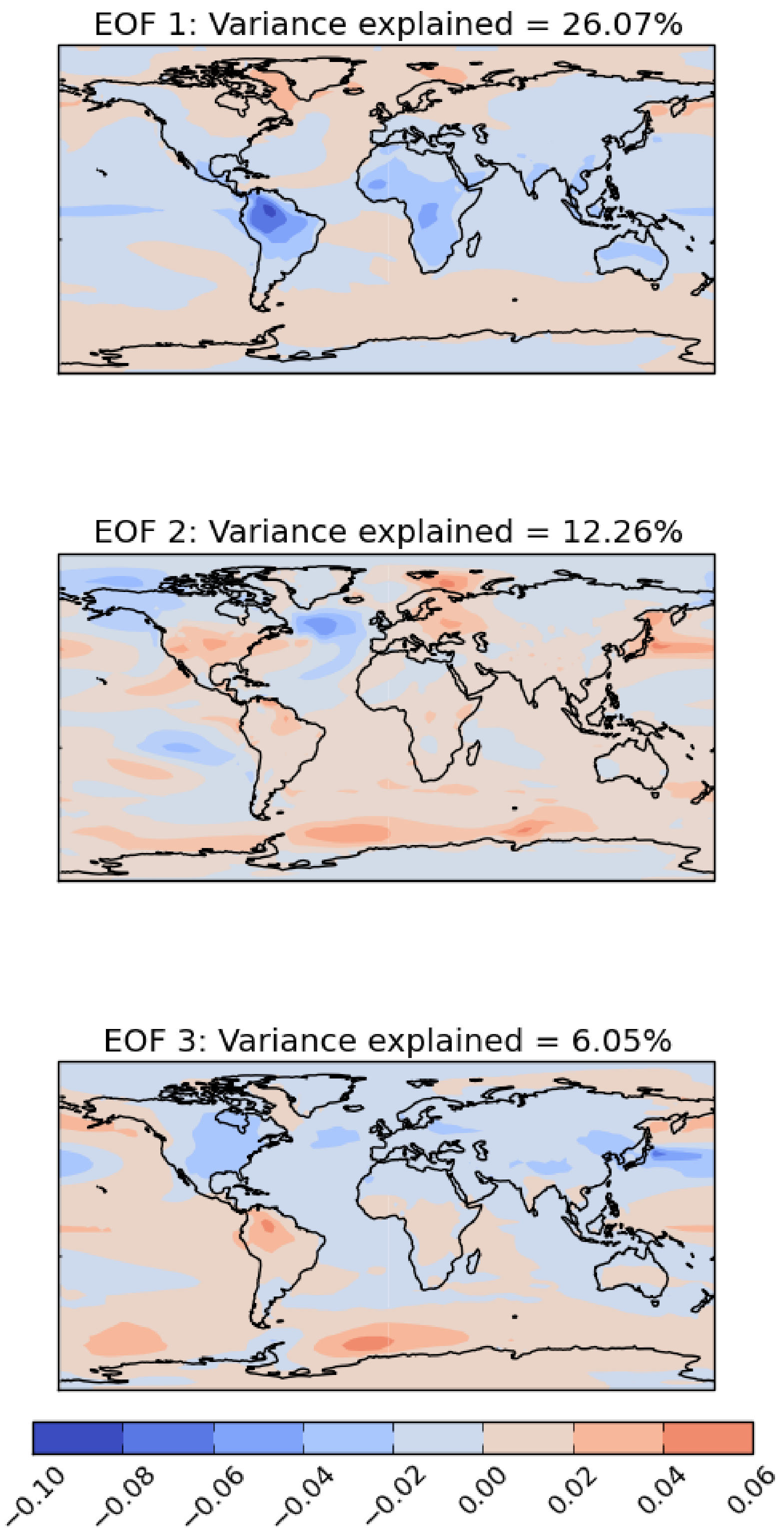


**Figure 3.9** – As for figure 3.8, but for the RAMP ensemble.

a set of perturbed parameter experiments that will be used as a sampling of diverse climate response to motivate further investigations into the limits of climate response uncertainty. The material in this chapter provided only a succinct discussion of how perturbations to the GCM sub-gridscale parameters influences the model physics. Whilst a full comprehensive description of the GCM and perturbed physics structure can only be gained by consulting the model technical documentation, it is hoped that this chapter offers the interested reader an introduction to the methods employed to attempt to sample uncertainties in GCM simulations of the physical mechanisms of climate response described in chapter 2.

Further to this, in this chapter I took a very large set of parameter combinations, chosen to have non-flux-adjusted and radiatively-balanced pre-industrial control climates, and conducted instantaneous doubling of $CO_2$ concentration and 1% annual concentration increase experiments. This chapter presented features of the climate response across the ensemble, to illustrate the diversity of climate response that is found within this sort of large PPE.

A wide variety of ECS and TCR is found within this ensemble when ECS is estimated through the x-axis intercept of a linear regression of net TOA energetic imbalance with the GMST anomaly. A long-tail of high ECS GCMs was found with ECS as high as 9.5K. However, very few models are found with ECS of less than 2K. The ensemble displays a great diversity in the simulated feedbacks and effective radiative forcing that combine to give the ECS. SW-CRE and LW-CRE feedbacks are found with both positive and negative signs at the global scale, as are positive and negative fast adjustments in SW-CRE fields. CRE forcing and feedback contribute the most to the range of ECS across the ensemble, with clear-sky forcing and feedback much more consistently simulated in different ensemble members. Across both global and regional quantities the instantaneous doubling experiments seem to be good guides to the physics represented in the time-evolving forcing scenario through linear models and local patterns of change show the same dominant modes of variation across the ensemble in both experiments.

The diverse representation of climate feedbacks within this ensemble shows the range of climate forcing and feedback mechanisms that can be explored within large PPE methods and arises from simple constraints on the control climate fidelity. In particular, it shows that these simple constraints are insufficient to rule out very high ECS and TCR, consistent with the results of other studies in the literature. In the next chapter the features of climate response from this ensemble will be compared with an alternative, observationally-based, method of estimating climate response uncertainty in order to help contrast the observationally-based ranges of climate response with ranges derived from GCM experiments.

# CHAPTER 4

## Energy-balance models of climate response

Chapter 3 showed a diverse range of climate response in a very large PPE of the HadCM3 GCM. However, capturing the spread in the properties of the GCM-simulated climate response for use in policy-relevant analyses requires computationally simpler models. This chapter focuses on energy-balance models of the climate system (section 4.1) and their use in both projections of future climate change (section 4.2) and for estimating the climate response from observations (section 4.3). This chapter then discusses the differences between GCM-based and observationally-based estimates of the climate response and proposes a novel metric for summarising the climate response, the use of which highlights a region of climate response space currently not sampled by GCMs but consistent with observationally-constrained energy-balance estimates (section 4.5). The majority of this chapter was published as a peer reviewed article in the journal *Climatic Change*.

## 4.1 Representing the climate response in simple climate models

Chapter 2 discussed how ECS is a function of the components of climate forcing and physical feedbacks associated with the $CO_2$-induced warming and how an *effective* ECS can be estimated from constant forcing simulations. Whilst for most plausible emissions trajectories the ECS is not the most relevant measure of the climate response, due to the long timescales and potential nonlinear feedbacks [Andrews *et al.,* 2012], ECS is still a highly relevant indicator of the physical climate response in concentration stabilisation scenarios and when considering the possibility of non-linear runaway feedbacks.

As discussed in chapter 2, the other canonical metric of the climate response is the TCR. Whilst not physically or statistically independent quantities, the ECS and TCR have conventionally been taken to be the standard metrics summarising the GMST climate response to $CO_2$. In order to relate the two metrics to each other, and to transient climate changes under more complex radiative forcing scenarios, some kind of climate model is always required. A variety of different choices can be made for the form and structure of a simple model of the climate system within the overall forcing-feedback framework [Gregory *et al.,* 2015].

### 4.1.1 The constant climate resistance approximation

One of the simplest model choices involves making the assumption that planetary heat uptake is linearly driven by GMST change, $\Delta N = \kappa \Delta T$ (known from here on as the "constant climate resistance" or CCR approximation). This is equivalent to assuming that the deep ocean acts as an infinite heat sink below an upper ocean of negligible heat capacity that loses heat to both space and the deep ocean proportional to $\Delta T$ as it warms. This allows planetary energy-balance (equation 2.1.1) to be rewritten as $\Delta F = (\kappa + \lambda)\Delta T$ [Gregory & Forster, 2008]. Assuming that the sum of $\kappa$ and $\lambda$ is constant over time, this model contains no response timescales and implies that GMST change is directly proportional to the radiative forcing change. Therefore, the TCR is derived from the simple model formulation as,

$$\mathrm{TCR} = \frac{F_{2\times}}{\kappa + \lambda} = \frac{F_{2\times}}{\rho}, \tag{4.1.1}$$

where $F_{2\times}$ is the effective radiative forcing from doubling atmospheric $CO_2$ concentrations. The utility of equation 4.1.1 depends on the constancy of the climate resistance, $\rho$ (where $\rho = \kappa + \lambda$), over time. Gregory & Forster [2008] show that this is a good approximation of GCM response for the historical period and for future radiative forcing scenarios, however this approximation cannot hold in concentration stabilisation scenarios, in which the system heat uptake eventually reduces to zero as the system approaches equilibrium.

### 4.1.2 The finite heat capacity upper-ocean model

Unlike the CCR model described in section 4.1.1, the real climate system does display a lag in the response of the climate system to a change in radiative forcing. Many different simple models can be invoked to describe this lag in the system. The next stage of additional complexity over the CCR approximation is to introduce a finite heat capacity to the upper ocean, whilst retaining the infinite heat capacity of the deep ocean. This simple model allows the net energetic imbalance of the climate system to be described as,

$$\Delta N = c_s \frac{\mathrm{d}\Delta T}{\mathrm{d}t} + \kappa \Delta T, \tag{4.1.2}$$

where $c_s$ is the finite (small) heat capacity of the upper mixed layer of the ocean.

Solving this equation under a constant radiative forcing, $F_{2\times}$, leads to a solution for the evolution of the surface temperature system,

$$\Delta T(t) = (\frac{F_{2\times}}{\rho})(1 - \exp[-\rho t/c_s]). \tag{4.1.3}$$

involving a single timeconstant of surface temperature response to the radiative forcing perturbation, $c_s/\rho$.

Integrated under a linear ramp in radiative forcing, the solution for TCR under this simple model formulation is given as,

$$\mathrm{TCR} = \frac{F_{2\times}}{\rho}\left(70 - \frac{c_s}{\rho}(1 - \exp[-\rho t/c_s])\right), \tag{4.1.4}$$

with the numerical factor of 70 arising from a doubling of pre-industrial occurring after 70 years when concentrations are increased by 1%/yr - see figure 2.5. Typically, this model formulation is most useful at short timescales when the upper ocean is not in equilibrium with the forcing perturbation.

### 4.1.3 The two-box impulse-response model

The real climate system does exchange heat between the upper mixed-layer of the ocean and the deep ocean, albeit slowly. The deep ocean has been observed to warm over the second half of the $20^{th}$ century [Levitus *et al.,* 2012] and many simulations of climate response to an abrupt change in $CO_2$ forcing require two time constants of response in order to account for the slow warming of the deep ocean [Held *et al.,* 2010; Geoffroy *et al.,* 2013]. Whilst the CCR approximation has been shown to work reasonably well for observed changes over the 20th century and for 1% per year $CO_2$ concentration increases [Raper *et al.,* 2002; Gregory & Forster, 2008], the real climate system displays lags in the GMST warming associated with the finite thermal inertia of the deep ocean.

An example of a simple climate model that includes two response timescales is an impulse-response model (or impulse-response function/IRF), an empirical representation of the GMST response to external forcing that has been tuned to emulate the GMST response of GCMs [Boucher & Reddy, 2008]. Mathematically, an impulse-response formulation is equivalent to the climate system's Green's function (the response of the system to a unit pulse, or $\delta$ function, of radiative forcing at time $t = 0$), similar to the step-response model discussed in chapter 3 for $CO_2$ forcing. Here, I use the impulse-response model of Boucher & Reddy [2008] to represent the time-dependent GMST change, $\Delta T$, due to a unit of radiative forcing at $t = 0$ as the sum of two decaying exponentials,

$$\Delta T(t) = \sum_i \frac{c_i}{d_i} \exp\left(\frac{-t}{d_i}\right), \qquad i \in [1, 2], \tag{4.1.5}$$

where $d_i$ is the characteristic timescale of each component of the climate response and the coefficient $c_i$ represents the relative magnitude of the timescale's contribution to equilibrium warming. The response timescales $d_1$ and $d_2$ are often taken to be physically consistent with the timescales of the atmosphere/land surface/upper ocean response and the slower response of the deep ocean respectively. Values can be estimated from long integrations of GCMs to be on the order of 10 and 400 years respectively [Boucher & Reddy, 2008]).

A two-time-constant impulse-response model is physically equivalent to a two-box climate model [Li & Jarvis, 2009]. Equation 4.1.5 represents the perturbation GMST, $\Delta T$, determined by a surface/deep ocean box model given by,

$$c_s \frac{\mathrm{d}\Delta T}{\mathrm{d}t} = F(t) - \lambda \Delta T - k(\Delta T - \Delta T_d) \tag{4.1.6}$$

$$c_d \frac{\mathrm{d}\Delta T_d}{\mathrm{d}t} = k(\Delta T - \Delta T_d), \tag{4.1.7}$$

where $\Delta T_d$ is the perturbation temperature of the deep ocean, $c_s$ and $c_d$ are the heat capacities of the surface layer and the deep ocean respectively, and $k$ is the heat exchange coefficient between the two layers.

The TCR and ECS can be expressed in terms of the impulse-response model parameters. The ECS under the impulse-response model is:

$$\text{ECS} = \int_0^\infty F_{2\times} \sum_i c_i/d_i \exp\left(-(\infty - t')/d_i\right) dt' = F_{2x}(c_1 + c_2). \tag{4.1.8}$$

TCR is the response to 1% annual $CO_2$ concentration increases, $C(t) = C_0(1.01)^t$, for 70 years to the time of double original atmospheric concentrations, which gives:

$$\text{TCR} = F_{2\times}\left(c_1\left[1 - \frac{d_1}{70}(1 - \exp\left(-70/d_1\right))\right] + c_2\left[1 - \frac{d_2}{70}(1 - \exp\left(-70/d_2\right))\right]\right). \tag{4.1.9}$$

Solving equations 4.1.6 and 4.1.7 for $\Delta T$ leads to a coupling between the timescales of the response $d_i$ and the magnitudes of each component $c_i$, when expressed in terms of the parameters of the two-box model. However, as the parameters of the two-box climate model are themselves uncertain, it is possible to consider the climate response functions consistent with various TCR and ECS combinations by varying the equilibrium response magnitudes, $c_i$, in each component whilst leaving fixed the response timescales, $d_i$.

Assuming known response timescales, equations 4.1.8 and 4.1.9 form a pair of linear simultaneous equations for the coefficients $c_1$ and $c_2$. This allows an impulse-response model to be constructed with TCR and ECS explicitly supplied as model parameters. A three-box formulation, including an additional sensitivity component and response timescale is required to represent the behaviour of GCMs over the longest time horizons [Li & Jarvis, 2009], but as the extra response timescale is on the order of a thousand years (and generally associated with timescales where additional Earth system feedbacks not included in GCM simulations, such as melting ice-sheets and biosphere changes [Lunt *et al.*, 2010], are important), it can be neglected for investigations of behaviour over the time horizon of the next centuries.

## 4.2 Implications of uncertainty in TCR and ECS for climate projections

The impulse-response model described in section 4.1.3 can be used to examine the impact of uncertainty in TCR and ECS respectively on projections of future climate change. Fixing the response timescales at values most representative of GCM simulations of the climate system ($d_1 = 8.4$ years and $d_2 = 409.5$ years, corresponding to HadCM3) [Boucher & Reddy, 2008], I vary the response coefficients, $c_1$ and $c_2$ to simulate the GMST climate response of any given combination of TCR and ECS. In this section this 2-box impulse-response model is used to examine the impact of changes in the IPCC uncertainty ranges for TCR and ECS between AR4 and AR5 on future projections of climate change.

### 4.2.1 Matching the impulse-response model to observations of present day GMST anomaly and radiative forcing

For a specific TCR and ECS, the IRF can be integrated with the observed historical radiative forcing followed by a future scenario of radiative forcing. However, to assess the consistency of different TCR and ECS impulse-response models with observations of historical climate change, uncertainty in historical radiative forcing must also be considered.

Here, I follow Padilla *et al.* [2011] in assigning all uncertainty throughout the historical radiative forcing timeseries to anthropogenic aerosols, motivated by the conclusions of current community assessments indicating that aerosol forcing is the largest component of contemporary uncertainty in globally averaged radiative forcing [Myhre, 2013a]. The timeseries of total ratiative forcing, $F(t)$, can therefore be expressed as,

$$F(t) = F_{other}(t) + xF_{aero}(t), \tag{4.2.1}$$

with $F_{aero}$ representing the best-estimate timeseries of the anthropogenic aerosol radiative forcing, $F_{other}$ as all other components of the radiative forcing, and $x$ as a non-dimensional scaling factor representing the aerosol forcing uncertainty. As the two-box IRF is linear in its forcing inputs, this scaling factor propagates directly into a scaling factor in the temperature response to the best-estimate anthropogenic aerosol radiative forcing, $T_{aero}$, with total GMST change given as,

$$T(t) = T_{other}(t) + xT_{aero}(t). \tag{4.2.2}$$

By constraining net temperature anomalies, as simulated by the 2-box model, to be consistent with observations over a defined period, it is possible to calculate the aerosol forcing scaling factor required to make a climate response function of specified ECS and TCR consistent with observations, whilst ensuring that total radiative forcing is also consistent with recent observational and modelled estimates. As an example, for a higher TCR climate response to yield an equal temperature change over an identical period, the historical radiative forcing over that period must have been correspondingly lower.

### 4.2.2 Projections of climate change under TCR and ECS uncertainty using the two box impulse-response function

To investigate the relative importance and impact of understanding TCR and ECS uncertainty for future projections of GMST change, the 2-box impulse-response model described in section 4.1.3 is integrated using the methodology of section 4.2.1 for a set of projections of possible global radiative forcing pathways for the coming century and beyond. These Representative Concentration Pathways (RCPs) are chosen from the integrated assessment modelling literature in order to span the range of conceivable radiative forcing evolution out to 2100 [Van Vuuren, 2011]. The number in the name of the four scenarios corresponds to the global anthropogenic radiative forcing achieved in 2100 (e.g. RCP8.5 has a net anthropogenic radiative forcing of 8.5$\mathrm{Wm}^{-2}$ in 2100), except for RCP-3PD which represents the peak anthropogenic radiative forcing before forcing declines to 2.6$\mathrm{Wm}^{-2}$ in 2100.

As an indicative sampling of climate response uncertainty I integrate the 2-box impulse-response function with upper/lower limits of the TCR and ECS *likely* ranges given in the Working Group 1 (WG1) reports in the IPCC AR4/5 respectively. For these TCR and ECS values the anthropogenic aerosol timeseries is scaled to give 2011 total radiative forcing consistent with the lower/upper *likely* estimates of 2011 total radiative forcing in AR5 respectively as well as agreement between the modelled GMST anomalies and the HadCRUT4 observational anomalies [Morice *et al.,* 2012] for the decade 2000-2009. The bounding values of these *likely* uncertainty ranges can be found in table 4.1. In AR4, there was no *likely* range for TCR given, with TCR uncertainty assessed as *very likely* greater than 1.0K and *very unlikely* greater than 3.0 K [Solomon, 2007], where *very likely* denotes a 9 in 10 probability and *very unlikely* a 1 in 10 probability. The only TCR *likely* range that can be inferred from AR4 without further assumptions is to set it equal to this interval. This is the most conservative estimate possible and avoids over-interpreting the result from AR4 whilst allowing comparison with the *likely* range from AR5. This means the lower limits of the *likely* ranges of TCR (1.0K) are taken as equal between AR4 and AR5 and likewise, the upper ECS *likely* range limit (4.5K) is taken as equal between AR4 and AR5.

Each of the four impulse-response model parameter sets are then integrated into the future under all four of the RCP pathways with their future anthropogenic aerosol radiative forcing subject to the same scaling factor required over the historical period. The time series of GMST anomalies are shown in figure 4.1.

| Parameter Set | TCR (K) | ECS (K) |
|---|---|---|
| AR4 - Lower *likely* | 1.0 | 2.0 |
| AR4 - Upper *likely* | 3.0 | 4.5 |
| AR5 - Lower *likely* | 1.0 | 1.5 |
| AR5 - Upper *likely* | 2.5 | 4.5 |

**Table 4.1** – The parameter values of TCR and ECS used to span assessed AR4/5 *likely* ranges of the physical climate response.

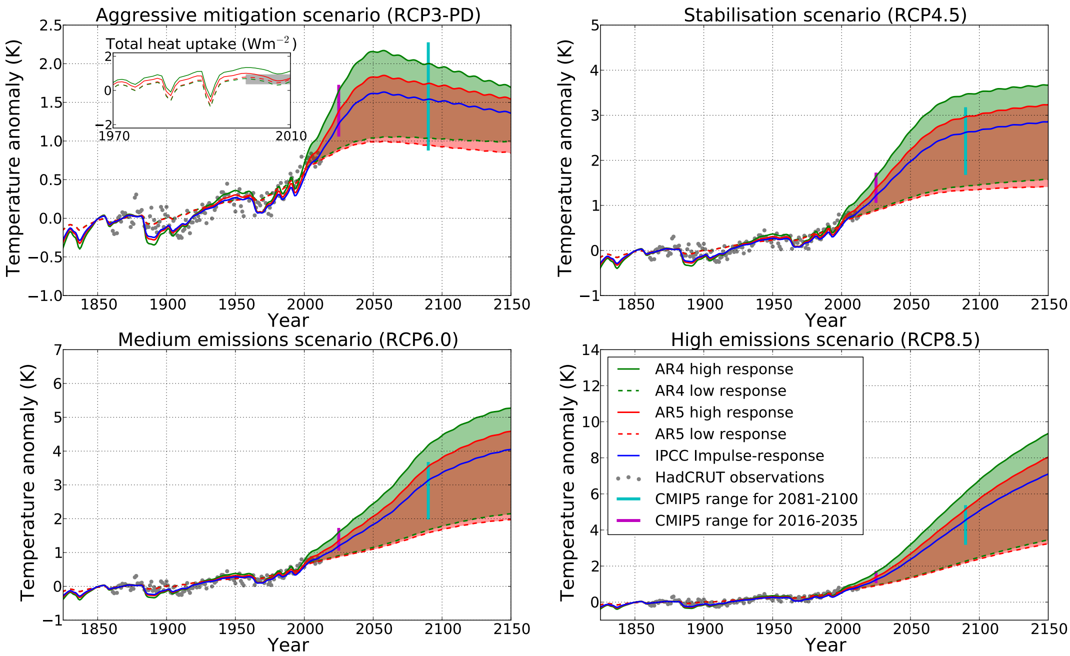


**Figure 4.1** – GMST projections for the four RCP scenarios derived using the impulse-response simple climate model. Note that the y-axis scales are not identical. The red shading shows the projection envelope consistent with AR5 ECS and TCR *likely* ranges and scaled forcings. The green shading shows the same for the AR4 *likely* ranges. All temperature changes are given relative to the 1861-1880 average. The range of the CMIP5 GCMs for the near-term and long-term GMST anomaly projections from AR5 are included. The blue line shows scenario projections under the impulse-response model used for multi-gas metric calculations in AR5 Ch 8 [Myhre, 2013a]. The inset panel shows system heat uptake for the limits of the *likely* ranges with the grey shaded bar marking the standard deviation around the 2000-2009 average.

The panels showing the projections for RCP6.0 and RCP8.5 in figure 4.1 show that the reduction of the lower ECS *likely* bound (the lower limits of the projection envelopes, in which the only difference between the AR4 and AR5 case is the lower ECS limit) in AR5-based projections has a barely discernible effect out until 2150, where the longer timescale climatic response, mainly associated with ECS, starts to emerge. The main effect of the AR5 TCR and ECS uncertainty updates is to reduce the likelihood of the high end of AR4-based projections, corresponding to the downgrading of the upper limit of the TCR estimate. Therefore, for policy-relevant projections, the revision of TCR range is much more important than revisions to the ECS range, as it reduces the upper edge of the projections envelope across the entire century whereas the ECS revision has a smaller effect on the lower edge of the envelope near the end of the century and beyond. An important caveat is that these temperature projections correspond to a climate system linearly responding to radiative forcing. Revisions to the ECS would be more important to future projections if significant non-linearities in the climate response to radiative forcing are encountered before the end of the century.

RCP3-PD shows a faster cooling after the temperature peak in the AR4 high-sensitivity case than the AR5 high-sensitivity case, despite a greater peak temperature change being achieved mid-century under the AR4 high-sensitivity integration. This is consistent with the higher TCR in the AR4 high-sensitivity limit causing a faster decrease of temperatures in response to net negative emissions technologies in the RCP3-PD scenario. However, it is important not to give too much weight to the projections of a simple climate model beyond the peak of any forcing timeseries [Schaller *et al.,* 2014], where issues of reversibility of climate response are present when forcing is falling.

It is clear that simple models which fit the observational record can still simulate a wide variety of climate futures. Unlike the long-term CMIP5 projections, the near-term CMIP5 projections span a range that extends above the projection envelope indicated by the AR5 TCR, ECS and forcing uncertainties. In all scenarios (except for RCP3-PD where the radiative forcing peaks and declines) the long-term GMST projections are either completely contained within the AR5 uncertainty envelope or extend beyond it by a lower percentage than in the near-term projections. This supports the scaling down of the raw CMIP5 GCM projections for the near-term warming (2016-2035) but an unadjusted long-term warming (2081-2100) when forming the *likely* projection ranges quoted in AR5 in order to be consistent with the sensitivity and forcing uncertainties quoted elsewhere in the report [IPCC, 2013]. In the RCP scenarios with fast growing radiative forcing (RCP6.0 and RCP8.5), the long-term CMIP5 projection ranges lie over the warmer part of the impulse-response model AR5 uncertainty envelope. This would be consistent with increases in the climate feedback parameter at higher GMST in the GCMs associated with potential additional feedbacks that are not captured in the fixed sensitivity simple model [Gregory & Forster, 2008; Winton *et al.,* 2010]. Using equations 4.1.8 and 4.1.9, the implicit TCR and ECS for the IRF used to calculate global temperature potentials (GTPs - see Fuglestvedt *et al.* [2010] for a discussion on metrics) in AR5 [Myhre, 2013b] can be shown to be TCR=2.21K and ECS=3.96K (the projections of the impulse-response model with these climate response parameters is marked as the solid blue line in figure 4.1). These numbers are on the higher end of the uncertainties for these parameters given in AR5, albeit still within the given *likely* ranges.

In this section, I have considered the literal interpretation of independent ECS and TCR ranges in AR4 and AR5 for future temperatures under standard emission scenarios. In allowing the TCR to change and ECS to be held constant between the high sensitivity limits of the AR4 and AR5 cases (and vice versa for the low sensitivity limits) figure 4.1 describes the relative impact of independent variations of the ECS and TCR on future projections. The extent that TCR and ECS could actually be considered independent will be a topic returned to later in this chapter.

## 4.3 The role of model structure in constraining climatic response using observations

Various methods have been used to constrain uncertainty about ECS and TCR. One approach is to use a simple model to constrain the joint distribution of TCR and ECS from observations of the planetary energy budget. Recent studies estimating TCR and ECS via these methods have received substantial attention [Otto *et al.,* 2013; Huntingford, 2013; Schmidt *et al.,* 2014], especially as they have suggested TCR uncertainty intervals that include lower values than sampled by existing GCMs. These methods, which incorporate direct observations of the climate system and do not use GCMs, are often referred to as observational estimates of TCR and ECS, as if they were independent of the simple models they use to interpret the observations. Many studies have focused on estimating ECS using an energy-balance model constrained by observations (e.g Forest *et al.* [2002], Aldrin *et al.* [2012] and Skeie *et al.* [2014]). It is well-established that inferences about ECS using such methods depend on the structure of the simple model used as shown by Knutti *et al.* [2008] in response to Schwartz [2007]. However, as TCR is more directly connected to observations of the planetary energy budget, observationally-based TCR estimates are often anecdotally referred to as 'model independent' without this assumption having been explicitly validated.

Constraining climate response from observations always requires some sort of a model to relate the observations to the climate response. This section contrasts the use of the two-timescale impulse-response model with the CCR approximation in order to show the impact of simple model structure (which I use to refer to both the mathematical formulation of the simple climate model and the value of parameters not varied when estimating climate response from observations) on inferences about the TCR. In this section I choose to focus on uncertainties in TCR estimates rather than in ECS estimates. This is both because structural dependence of uncertainties in observationally-based estimates of ECS are already well-established [Knutti *et al.,* 2008], and because the varying radiative forcing in the TCR definition is more relevant to policy decisions that have to be taken over the next few decades in which the radiative forcing is likely to still be increasing and the climate system is substantially out of equilibrium due to heat uptake by the oceans. Indeed, the analysis in section 4.2 shows that when TCR and ECS are considered independent the TCR is more important for climate projections out until around 2150. Again, noting caveats given before about very high ECS and runaway greenhouse effects, TCR is likely to prove a more relevant metric of climate response to future projections as it better predicts the transient response to cumulative emissions (TCRE) [Gillett *et al.,* 2013; Allen & Stocker, 2014].

### 4.3.1 Observations of the planetary energy budget

In order to investigate the role of model structure in observationally-based estimates of climate response a set of observations of the planetary energy budget and associated uncertainties are required. Recent studies have focused on estimating TCR and ECS from observations of the most recent complete decade, 2000-2009 [Lewis, 2013; Otto *et al.,* 2013; Lewis & Curry, 2014], which has received particular interest due to the so-called 'hiatus' in the GMST warming (top panel of figure 4.2 in which the linear trend across 2000-2009 is merely 0.05K/decade). There has been significant interest in this decade and in particular whether the existence of such a hiatus, despite the continued increase in net forcing on the system, is evidence for lower ECS and TCR than was previously assumed (e.g. Huber & Knutti [2014]), with several energy-balance based studies suggesting that the implied sensitivity ranges are lower than previous community assessments and lower than as simulated in GCMs [Otto *et al.,* 2013; Lewis & Curry, 2014].

To sample uncertainties in the planetary energy budget over this decade, I construct a large number of realisations of energy budget components consistent with the observations. For GMST anomalies, the HadCRUT4 dataset [Morice *et al.,* 2012] is comprised of a 100 member ensemble, that systematically samples uncertainties in the method of constructing a single timeseries from the raw station records. The members of the ensemble are shown in the top panel of figure 4.2. As each realisation is a separate timeseries, auto-correlations in the data, associated with multi-year internal variability, is correctly captured and uncertainty can be assessed by via the spread of the ensemble. The ensemble members form an approximately Gaussian distribution of the decadal GMST anomaly over 2000-2009 from which the standard deviations of the Gaussian can be computed. The ensemble mean anomaly over 2000-2009 is 0.75K and the standard deviation is 0.09K. Following Otto *et al.* [2013] I also include an additional error covariance with standard deviation of 0.08K added in quadrature to the raw HadCRUT4 standard deviation to account for internal variability around the true signal in the observed timeseries. This estimate is derived from the upper end of the CMIP5 internal variability range. This gives an overall standard deviation in the GMST anomaly for 2000-2009 of 0.12K.

To construct an ensemble of global-mean radiative forcing, I decompose the total radiative forcing into the 11 components assessed separately in Myhre [2013a]. Myhre [2013a] gives best-estimates and 5-95% confidence intervals for the 2011 radiative forcing. The volcanic forcing uncertainty has a small value in 2011 and hence uncertainties are capped at $\pm$ 30% to avoid an over estimate of the uncertainty of the magnitude of volcanic forcing after the Mt Pinatubo eruption in 1991. Based on these confidence intervals, I assume that the uncertainty is asymmetrically Gaussian distributed (a different standard deviation consistent with the asymmetric confidence interval either side of the best-estimate) allowing an ensemble of plausible values for the components of 2011 radiative forcing to be constructed with draws from a standard normal distribution. The best-estimate timeseries of each component is then scaled by the fractional change in the 2011 uncertainty for each draw to give a realisation for the entire radiative forcing history of that component. The ensemble of total radiative forcing timeseries is shown in the middle panel of figure 4.2. Using this method of draws from a Gaussian distribution has the advantage of having an equal number of draws from below the best-estimate as above it. Therefore, the confidence intervals from resulting distributions in TCR and ECS can be evaluated from the corresponding percentiles of the distribution and the best-estimate found by locating the median of the distribution.

For the ensemble of total planetary heat uptake, I exactly follow the method of Otto *et al.* [2013]. The ocean heat uptake is responsible for the vast majority of the increase in the planetary heat content with observational data for the upper ocean heat content derived from Domingues *et al.* [2008] updated to 2012 [Church *et al.,* 2013] and for the deep ocean (700-2000m) from Levitus *et al.* [2012]. Added to this are estimates of the abyssal oceanic heat content and contributions to the planetary heat content from the melting of ice and energy stored in the atmosphere as in Otto *et al.* [2013]. From this timeseries of planetary heat uptake, the rate of decadal heat uptake change can be calculated for 2000-2009. This is expressed as an anomaly relative to the period 1860-1879 as done in Gregory *et al.* [2002].

### 4.3.2 Method of estimating distributions of TCR and ECS

A model is required to propagate observational uncertainty into values of climate response parameters. As the two-box model is a linear combination of the long and short response terms, it is possible to express the model temperature anomaly for the 2000-2009 period, $\Delta T_m$, as,

$$\Delta T_m = c_1 b_1^* + c_2 b_2^*, \tag{4.3.1}$$

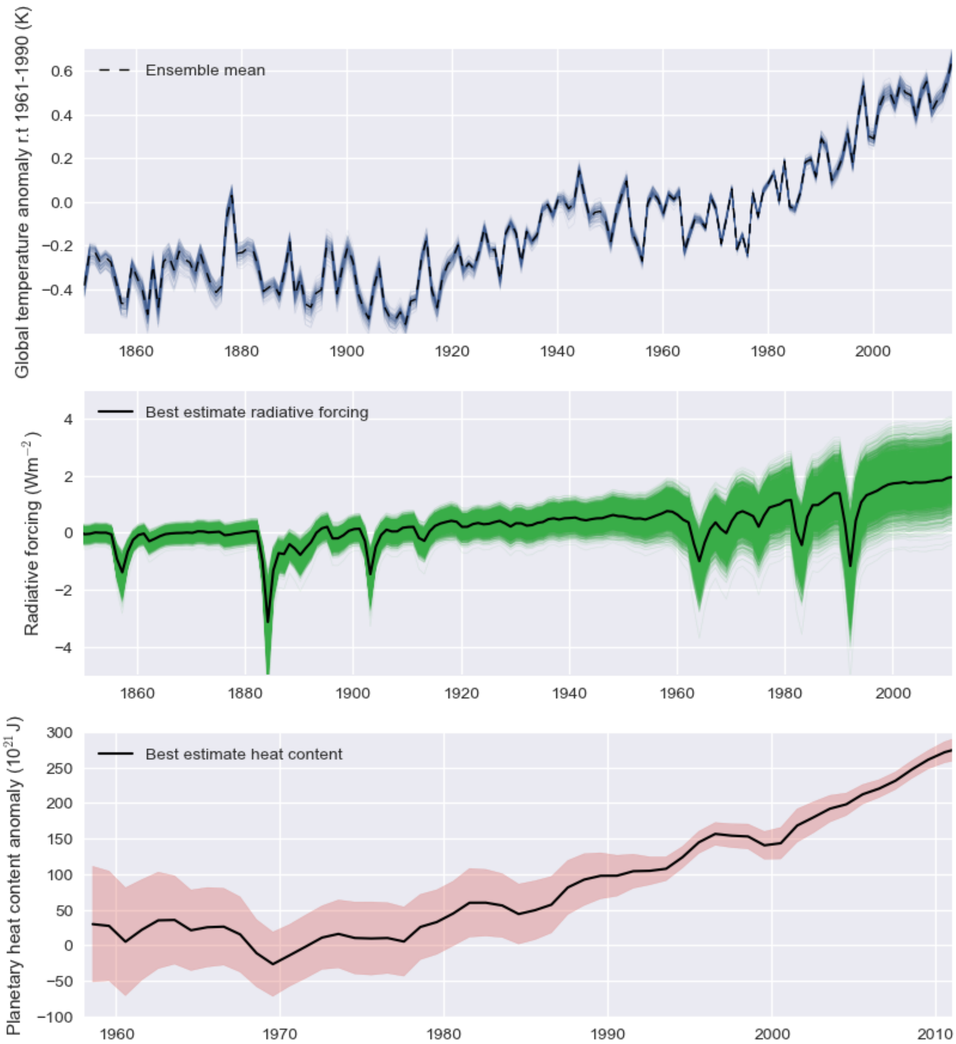


**Figure 4.2** – Observational data on the planetary energy budget. The top panel shows GMST anomalies (expressed relative to the 1961-90 period). The blue lines are individual members of the HadCRUT ensemble, with the dashed black line representing the ensemble mean. The middle panel shows the ensemble of radiative forcing uncertainty (green lines), generated by multiplicative scaling of the forcing component timeseries by uncertainty in their 2011 values. The solid black line shows the best-estimate for the total radiative forcing on the system. The bottom panel shows the total system heat content anomaly (black line) relative to 1961 with red shading either side indicating one standard deviation uncertainty.

where, $c_i$ are the model sensitivity coefficients and $b_i$ is the normalised temperature response of the short/long time scale model term respectively, given as,

$$b_i(t) = \int_{1750}^{2009} \frac{F'(t')}{d_i} \exp\left[-(t - t')/d_i\right] dt', \tag{4.3.2}$$

with $F'(t)$ a specific realisation of the radiative forcing ensemble and the $^*$ indicating an averaging over 2000-2009 and expression as an anomaly relative to 1860-1879. Computing the integrated responses of the model, $b_i^*$ (assuming fixed response timescales) allows the simple model parameters, $c_1$ and $c_2$, to be expressed as,

$$c_1 = \frac{\Delta T' b_2^*}{(\Delta F' - \Delta N' - \Delta T')(b_2^* - b_1^*)} \tag{4.3.3}$$

$$c_2 = \frac{\Delta T' - c_1 b_1^*}{b_2^*}, \tag{4.3.4}$$

where $\Delta F'$, $\Delta N'$ and $\Delta T'$ are specific realisations of the observational ensembles of radiative forcing, planetary heat uptake anomaly and GMST anomaly for 2000-2009 respectively.

From this distribution of simple model parameters, the set of simultaneous equations in equations 4.1.8 and 4.1.9 can be used to calculate the corresponding distributions of TCR and ECS. To this distribution I apply a physicality constraint that TCR and ECS must both be positive. Equations 4.1.8 and 4.1.9 require a value for the the radiative forcing of doubling $CO_2$, $F_{2\times}$ to estimate TCR and ECS. The forcing ensemble was derived assuming a best-estimate $F_{2\times}$ of 3.74 with a $\pm 20\%$ standard deviation uncertainty. This uncertainty is associated with many factors including the fast adjustments of the troposphere that remain difficult to constrain [Andrews *et al.*, 2012]. The correlation of $F_{2\times}$ with the other terms of the forcing distribution, except for the $CO_2$ timeseries, is unclear. It seems conceivable that non-$CO_2$ forcings could be positively correlated, negatively correlated or uncorrelated with the magnitude of the $CO_2$ radiative forcing, making the correlation between the value of $F_{2\times}$ and the net historical forcing unclear. Therefore, for simplicity and clarity, I choose to assume that the value of $F_{2\times}$ used to calculate TCR and ECS distributions is correlated with the uncertainty on the 2011 $CO_2$ forcing whilst leaving all the uncertainties of the different components of the timeseries unadjusted and uncorrelated with the $F_{2\times}$ uncertainty.

An alternative method to the sampling of observational uncertainties used here would be to sample the coefficients $c_1$ and $c_2$ explicitly, or directly sampling the underlying physical simple model parameters. However, this would lead to complex issues over what particular sampling distributions of model parameters would imply for distributions of the observations, as the observations are not simple linear functions of these parameters. Here I choose to leave this approach to future studies. In particular, these questions have been explored as part of a masters project arising from the work in this chapter.

### 4.3.3 The sensitivity of observationally-based estimates of TCR and ECS to model structure

The CCR approximation, in which temperatures respond instantaneously to the forcing, is the limiting case of the impulse-response model in which the short time-constant tends to zero and the long time-constant tends to infinity. This makes the approximation for TCR in equation 4.1.1 exact. Under a linear forcing ramp in which $CO_2$ concentrations are increased by 1% each year for 70 years, the CCR approximation underestimates the GCM-calibrated impulse-response model's actual TCR due to the inclusion of thermal delay in the impulse-response model (not shown).

The reducibility of the two-box impulse-response model to the single timescale climate response function and the CCR approximation allows the dependence of TCR confidence intervals on model structure to be examined by altering the value of the fixed response timescales of the simple climate model.

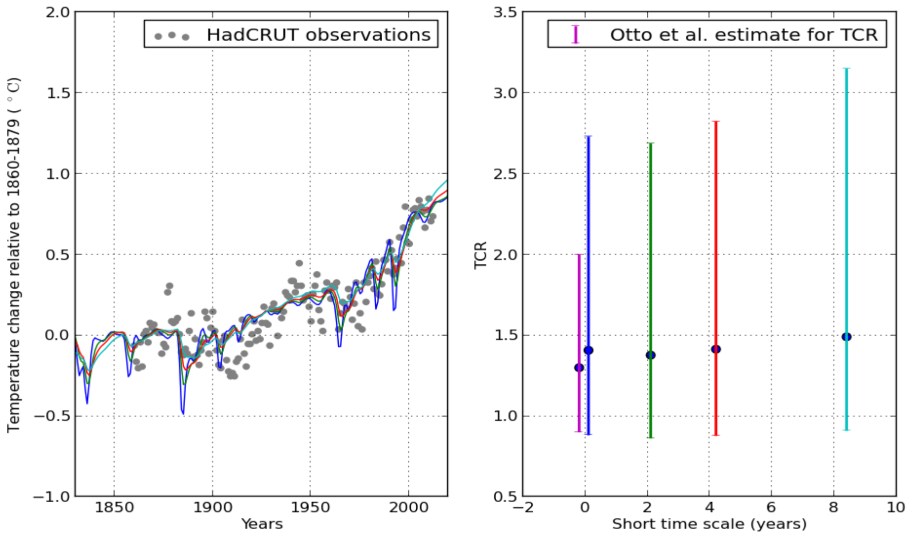


**Figure 4.3** – The left-hand panel shows the best fit model to observations with a variety of values for $d_1$, the short timescale in the impulse-response model. The right-hand panel shows the best-estimate and 5 - 95% confidence intervals for TCR for different values of $d_1$. The purple bar shows the estimate derived in Otto *et al.* [2013]. Its position on the x-axis is arbitrary.

The panels in figure 4.3 show inferences about TCR and its uncertainty from the observational data as a function of the short timescale of the climate response used in the IRF, $d_1$. As the short response timescale is the dominant mode over the observed historical period, inferences about TCR are largely insensitive to the long response timescale (not shown).

The observational GMST data can be fit well for a variety of impulse-response models (the left-hand panel in figure 4.3). The main difference of fit between the models occurs where they are responding to strongly negative but short-lived forcing from volcanic eruptions. These appear (assuming all forcing efficacies to be unity) to allow us to exclude the shortest timescale version of the impulse-response (equivalent to the CCR model) but do not effectively distinguish the other models. Using volcanoes to constrain response timescales is also complicated by dynamical features of the volcanic response not captured by an energy-balance model.

The right-hand panel of figure 4.3 shows that the choice of short timescale leads to considerable difference in the upper bound of the TCR confidence interval. I here focus on frequentist confidence intervals to make my approach more transparent and avoid issues over prior distributions (see Lewis [2013]). I show distributions for multiple values of the short response timescale solely in order to demonstrate structural uncertainty in this method. As I am not attempting to construct a Bayesian posterior representing a degree of belief in the exact TCR of the real climate system, I do not comment on the relative likelihoods of different short response timescales and the discussion here serves merely to highlight the sensitivity of the results to this parameter. While the best-estimate TCR and lower confidence interval bound are relatively insensitive to $d_1$, the upper confidence interval bound shows a non-monotonic dependence on $d_1$, demonstrating that the choice of simple climate model is important. The purple bar shows the 5-95% confidence interval on TCR from Otto *et al.* [2013]. That study used the CCR approximation, with the deviation from the blue bar (the CCR model considered in this chapter) arising from the updated radiative forcing uncertainties used here. The increased radiative forcing uncertainty again primarily impacts the upper bound of the TCR confidence interval. In particular the increased magnitude of the aerosol forcing in the global-mean has the effect of making higher TCR values more consistent with the observed record.

This section showed a clear dependence of the 95% TCR confidence intervals derived from energy-balance data on the structure of the simple model. Whilst the 5% TCR confidence interval derived via these methods is insensitive to the fast-response timescale, in probabilistic risk assessments of climate impacts and emission mitigation polices it is often the upper confidence interval is most important (see chapter 1). It is clear that there is no such thing as a model-independent estimate of the climate response from observationally-based energy-balance methods.

Winton *et al.* [2010] reframed the two box model of the climate system (equations 4.1.6 and 4.1.7) to introduce a efficacy factor, $\epsilon$, on the radiative effect of a unit of ocean heat uptake. This efficacy factor physically corresponds to differing radiative effects of different patterns of ocean heat uptake of the same global magnitude

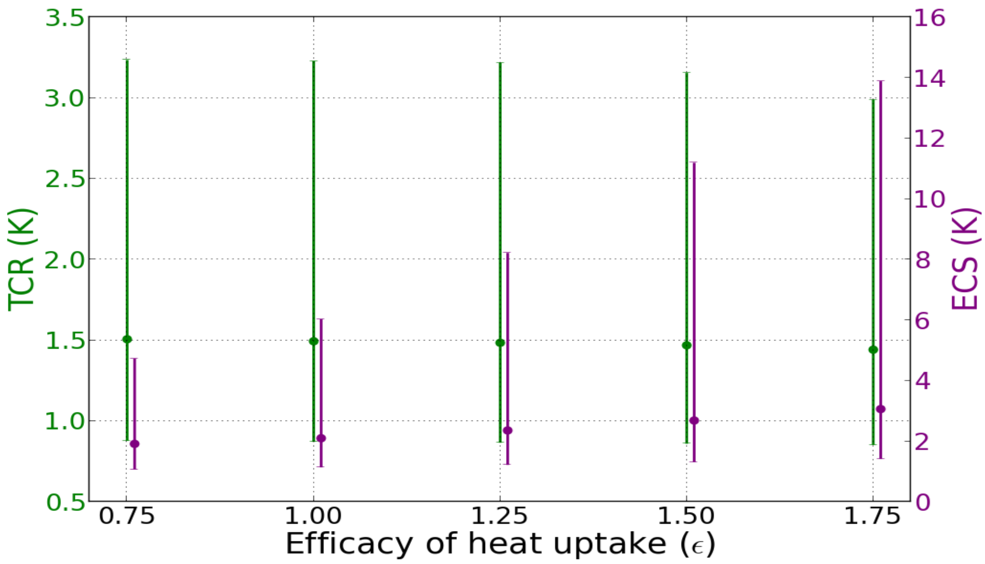


**Figure 4.4** – Sensitivities of the 5-95% confidence intervals to the value of the efficacy of ocean heat uptake. The left axis shows estimates for TCR (green) and the right axis for ECS (purple). All intervals are estimated for the 2000-2009 period using $d_1 = 8.4$ years.

(when the oceanic heat uptake is treated as equivalent to a negative radiative forcing, see section 2.2). If $\epsilon > 1$, the heat uptake that is observed in the climate system has a stronger radiative effect per unit energy flux than $CO_2$, suppressing radiative response of the surface climate system (and therefore warming) that would later become apparent as the heat uptake reduces to zero when the system approaches equilibrium. The possibility of $\epsilon > 1$ can therefore have important impacts on any estimates of the ECS (as shown in figure 4.4). Increases in $\epsilon$ from 1 to 1.5 (CMIP5 GCMs typically have $\epsilon$ between 1 and 1.5 at the time of concentration doubling in the 1%/yr concentration increase experiments [Winton *et al.*, 2010]) almost doubles the upper ECS confidence interval. In contrast, the upper confidence interval for TCR is relatively unaffected by the value of $\epsilon$ used, as would be expected due to the closer relation of the TCR to the observed climate state of today. Figure 4.4 shows confidence intervals for the GCM-calibrated impulse-response function ($d_1$=8.4 years), but dependencies on $\epsilon$ are largely insensitive to the value of $d_1$ (not shown). As $\epsilon$ is very difficult to constrain without GCM simulations, uncertainties in ECS from observationally-constrained energy-balance estimates are inherently sensitive to the value assumed, questioning the utility of attempting to constrain ECS from observations. As TCR is not similarly dependent on this assumption, future energy-balance studies should focus on estimating uncertainties in this more robustly observationally-constrained metric of climate response.

## 4.4 Comparing climate response in GCMs and observations

An obvious question that arises out of the results of section 4.3.3 is how these estimates, derived from observations using a simple energy-balance-model, compare to GCM simulations over the same period? In particular, is it the case that the GCMs can be shown to be inconsistent with the confidence intervals over the last decade and if so does that indicate an inconsistent GCM-simulated TCR with the observationally-based estimates?

When comparing climate models with observations, it is necessary to account for the observational mask (the finite locations over which observations exist at different times). Jones *et al.* [2013] consider the masking of the CMIP5 GCMs in attributing observed GMST changes to anthropogenic and natural factors. In this section, I use the historical CMIP5 simulations (extended with RCP4.5 forcings after 2005) from Jones *et al.* [2013] with the HadCRUT4 observational data mask applied to the simulations in order to calculate the 2000-2009 averages of the temperature anomaly masked to consistent locations with the observational data. Radiative forcing across the historical simulation is derived from Forster *et al.* [2013] in which the climate feedback parameter is calculated from a instantaneous $CO_2$ quadrupling experiment, and is then assumed to be the same for different forcing agents in order to calculate a timeseries of radiative forcing over the historical period from simulated GMST and heat uptake responses. The effective radiative forcings are given as global-mean quantities and are not masked with the Had-

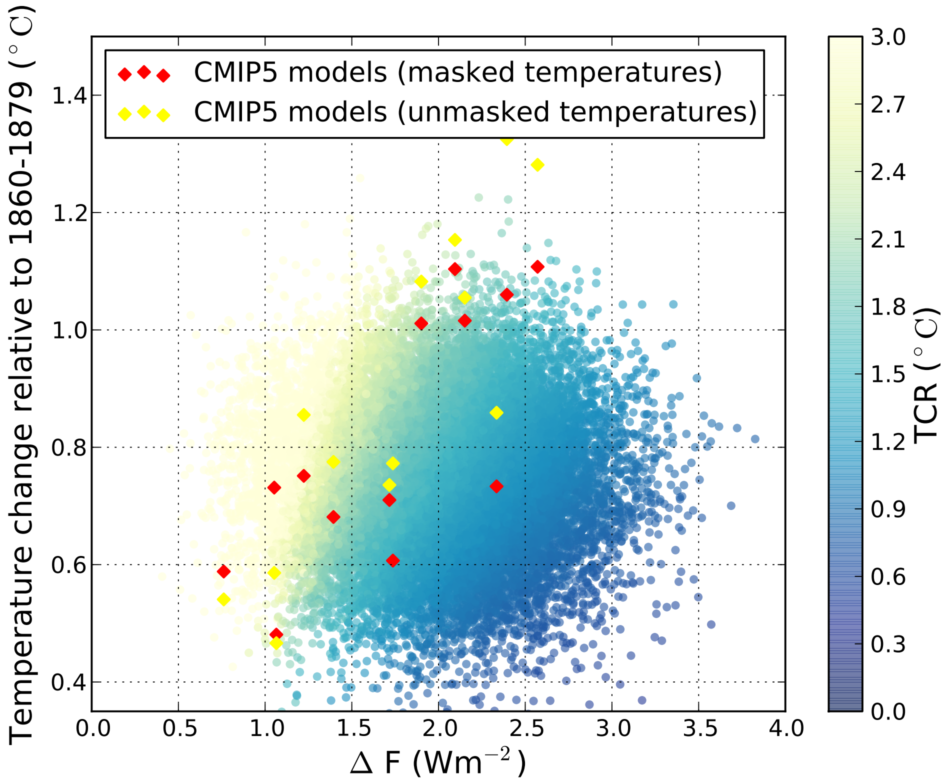


**Figure 4.5** – Comparing GCM historical energy-balances and observations. The x-axis shows the radiative forcing averaged over 2000-2009 with the y-axis showing changes in the GMST between 2000-2009 and the 1860-1879 reference period. The yellow diamonds represent CMIP5 GCMs using GMST anomaly values calculated only over the grid points covered by the observational mask. The red diamonds show the same GCMs with GMST anomalies calculated using the full unmasked spatial GMST field. The CMIP5 GCMs shown are those with available data for both the masked and unmasked GMST anomalies as well as the historical radiative forcing. The coloured circular dots represent individual samplings of observational uncertainty where the colouring is by the TCR as estimated using the two-box impulse-response model.

CRUT4 observational mask.

Figure 4.5 shows individual realisations from the ensemble of observational uncertainty in both temperature change and radiative forcing over the 2000-2009 period coloured by the TCR value for that member when calculated using the GCM-calibrated version of the two-box impulse-response model as in section 4.3.3. Diamonds show the masked and unmasked values for the CMIP5 GCMs for which forcing data, masked GMST anomalies and unmasked GMST anomalies were all available. Whilst the exact effect and magnitude of the masking differs between

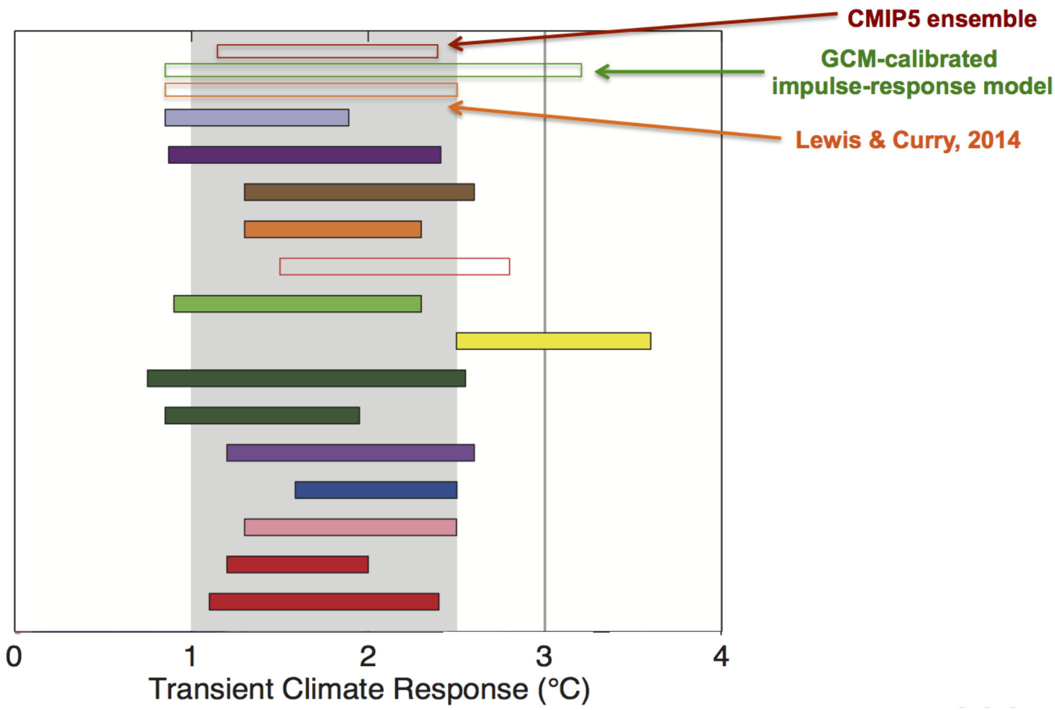


**Figure 4.6** – Figure adapted from IPCC AR5 TFE.6, Figure 2, showing 5-95% confidence intervals for TCR for various observationally-based studies. Added to this are the intervals from this chapter (green), the Lewis & Curry [2014] study (orange) and the CMIP5 range (reddy brown), all labelled with arrows at the top of the figure. The grey shaded bar represents the assessed *likely* range for TCR in IPCC AR5.

GCMs, the general effect of accounting for the observational mask is to reduce the temperature anomaly over the 2000-2009 period. This effect can be explained largely due to the absence of surface-based observations in the Arctic, which has been warming faster than the rest of the globe under Arctic amplification [Cowtan & Way, 2014].

Whilst not all CMIP5 GCMs shown are inconsistent with the ensemble of observationally derived 2000-2009 anomalies shown here, the effect of masking the GCM temperature output serves to, in general, push the CMIP5 GCMs closer to the centre of the observational distribution. The computed TCR values for the observational ensemble (colouring of the circular dots in figure 4.5) shows that by not accounting for the observational mask these observationally-based methods can underestimate the true global TCR for CMIP5 GCMs.

The effect of observational masks is one reason why the raw GCM global TCR values might be expected to be higher than observational estimates. Other limitations of these methods, including the varying efficacies of different forcing agents and the non-linearity of the forcing-feedback regime are discussed in section 4.6. Even accounting for the likely low bias of this method, the results from section 4.3.3 demonstrate that accounting for simple model structural effects in these observationally-based energy-balance estimates would be more than sufficient to span the GCM range (see figure 4.6). The constancy of the lower bound of TCR uncertainty under changes to the simple model structure does, however, highlight a region of low TCR that is potentially consistent with the recent observations but not within the CMIP5 GCM range. This is discussed further in the next section.

## 4.5 Climate response space in GCMs

This section investigates the relationship between TCR and ECS in GCMs and introduces the concept of realised warming fraction (RWF) as an additional metric of climate response.

### 4.5.1 The Realised Warming Fraction

The TCR and ECS are both measures of the climate response to external radiative forcing, but on different time-scales. In GCMs ECS is approximately linearly related to TCR over low to mid range realisations of ECS [Knutti *et al.,* 2005; Forster *et al.,* 2013]. Figure 4.7 shows this relationship in both the CMIP5 MME and the large and diverse radiative-balanced climate response PPE of the HadCM3 GCM described in chapter 3.

Figure 4.7 includes the IPCC *likely* uncertainty ranges for ECS and TCR from the IPCC-AR4 (green) and IPCC-AR5 (red) [IPCC, 2007; IPCC, 2013]. IPCC estimates for TCR and ECS are given independently without reflection of the correlations shown in figure 4.7. These independent ranges do not explicitly rule

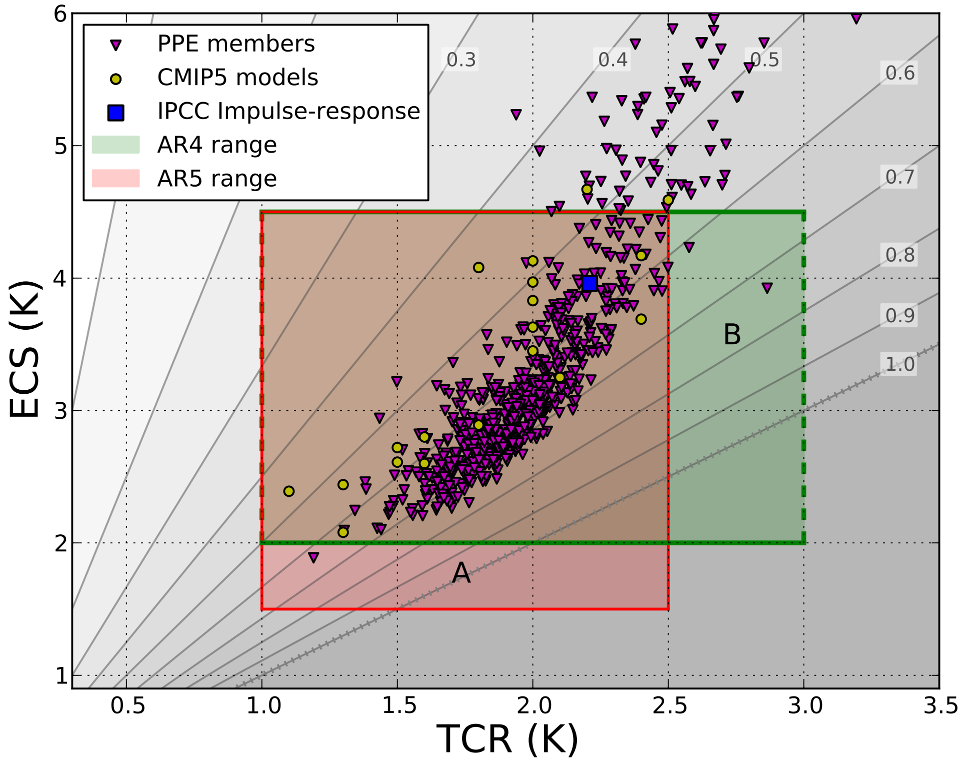


**Figure 4.7** – TCR and ECS for the CMIP5 MME (yellow circles) and for the PPE of the HadCM3 GCM (purple triangles). The *likely* ranges from AR5 (red) and inferred AR4 (green) are shown with grey lines denoting constant RWF (TCR:ECS ratio). The impulse-response model used for multi-gas metric calculations in AR5 is marked as "IPCC Impulse-response". Regions not common to both AR4 and AR5 *likely* estimates are marked A and B.

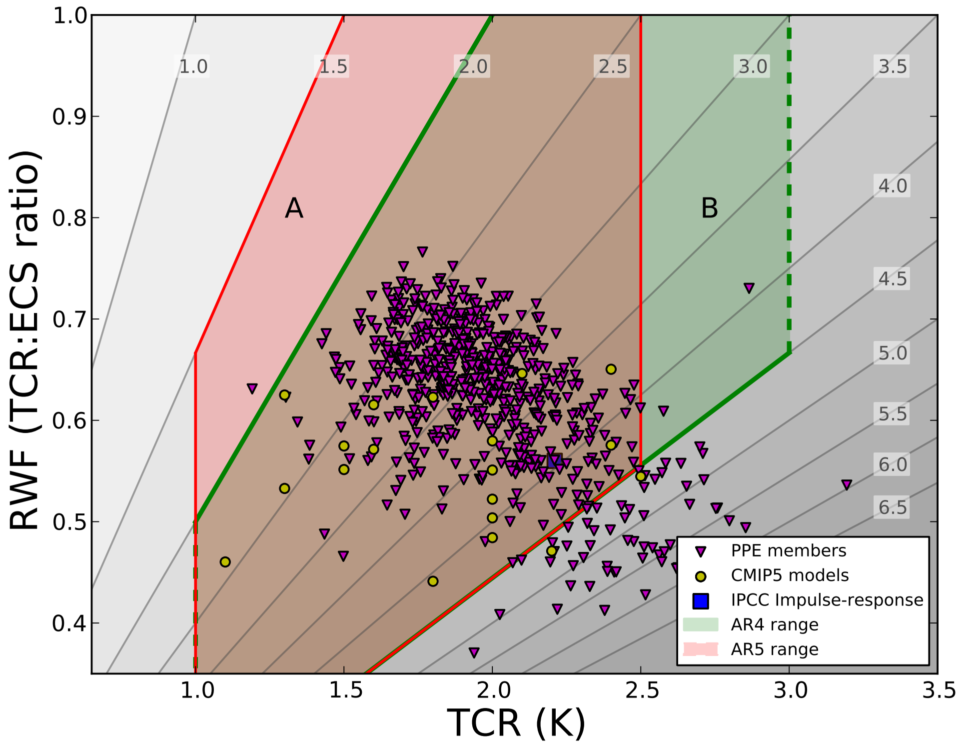


**Figure 4.8** – As figure 4.7 but the y-axis is the RWF (TCR:ECS ratio) and now lines of constant ECS are marked and shaded in grey.

out regions of the TCR-ECS joint distribution that that have TCR>ECS. Taking this additional constraint (the TCR = ECS line is represented by the dashed grey line in figure 4.7) would rule out parts of region A and B in figure 4.7 as physically implausible.

Figure 4.8 shows an alternative way to visualise the information shown in figure 4.7. The TCR:ECS ratio used as the y-axis in figure 4.8 is a measure of the fraction of committed warming (the equilibrium warming for a fixed concentration change) already realised after a steady increase in radiative forcing, a clearly policy-relevant quantity. I call this quantity the *realised warming fraction* (RWF) hereafter. The lower correlation between the RWF and the TCR demonstrates that these two quantities are more statistically independent in contemporary GCMs than TCR and ECS (correlation coefficients of 0.12 and 0.86 respectively for the CMIP5 GCMs). Whilst the RWF and TCR are mathematically linked (via ECS), this approximate statistical independence in state-of-the-art GCMs makes characterising the climate response in terms of these quantities more informative than in terms of the more correlated TCR and ECS. There is evidence of negative correlation between RWF and TCR in the PPE (particularly at high TCR). This is consistent with the frequently noted deviation of the TCR vs. ECS relationship from linearity at high climate sensitivities [Knutti *et al.,* 2005], which are sampled more densely in the PPE than the MME.

It may be easier to understand the implications of climate response revisions for climate projections if studies focus on estimating these two near-independent quantities, which can both be estimated from observations (RWF can be estimated from observations via a simple model, e.g. equation 4.5.1), although estimates of any two are of course sufficient to determine the third. It could be more useful to policy-makers for a body such as the IPCC to report uncertainty ranges in these two more independent quantities rather than independent ranges for the correlated TCR and ECS. A choice of more independent parameters would communicate more information about the climate response. Andrews & Allen [2008] discusses how independent metrics of climate response are preferable for understanding climate response within a simple model framework. The use of RWF and TCR as independent metrics of climate response fulfil this criteria and is equally applicable to simple models and GCMs.

Figure 4.9 shows the 90% likelihood regions for the GCM-calibrated impulse-response model and the CCR approximation in the two-dimensional space of figure 4.8 and figure 4.7 respectively. The dependence of the inferences about the climate response on simple-model structure is shown via the difference between the two contours. However, a common feature of the inferences from observations using the simple models considered here is a region of high RWF at low TCR consistent with these observationally-based methods but not sampled by the GCMs. The HadCM3 PPE samples higher RWFs than the CMIP5 ensemble (although not at the same time reaching the low TCR values necessary to match the contours from the simple models). However, these GCMs have had no tuning or validation beyond

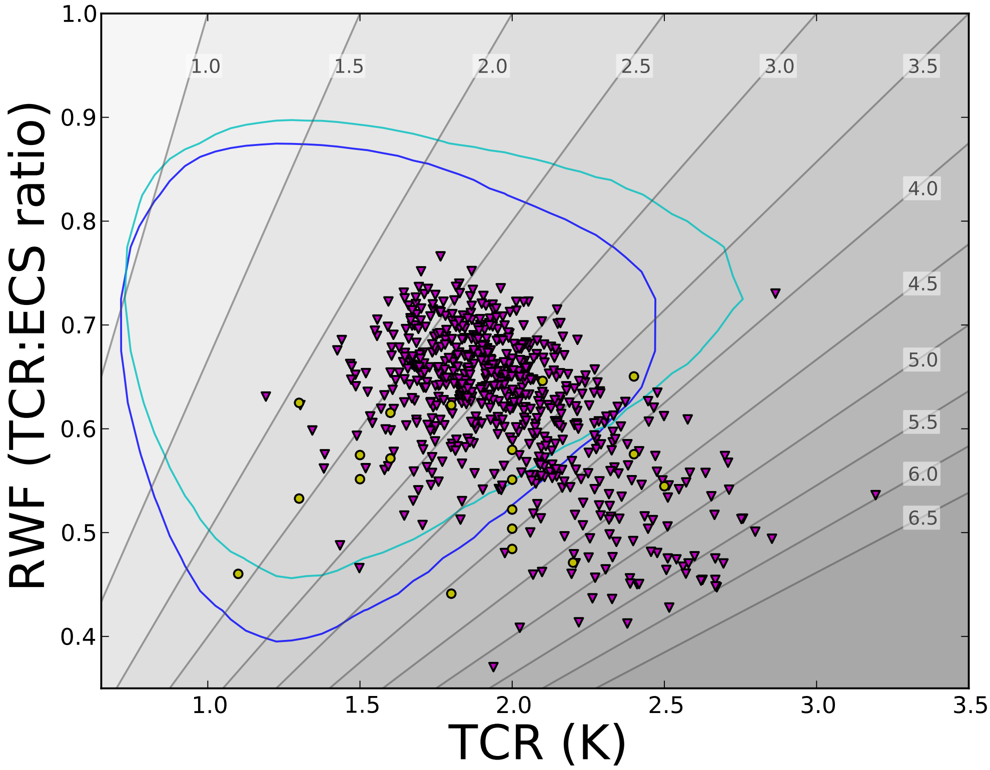


**Figure 4.9** – The 90% maximum likelihood regions for the GCM-tuned impulse-response model (light blue) and CCR approximation (royal blue) on the same axes as in figure 4.8.

their ability to reproduce TOA fluxes in a control climate (chapter 3).

Neither the HadCM3 PPE nor the CMIP5 MME give an indication of a decrease in RWF for low TCR. Similarly, recent observational estimates of TCR and ECS using simple climate models are also consistent with a high RWF at low TCR (e.g. [Otto *et al.,* 2013]). As such, the AR5 *likely* region, which includes region A in figure 4.8 whilst excluding region B, represents a region more consistent with the overall understanding of TCR and ECS uncertainties based on published evidence from multiple methods.

Under the CCR approximation, the RWF can be written as,

$$RWF = \frac{\widehat{TCR}}{\widehat{ECS}} = 1 - \frac{\kappa}{F_{2\times}}\widehat{TCR} = 1 - \frac{\Delta Q}{\Delta F}. \tag{4.5.1}$$

Here $\hat{x}$ is used to denote an observational-estimate for a property of the system, $x$.

$\widehat{TCR}$, might be called an effective TCR (as it is estimated from observations), but in many ways this would be misleading as, unlike the climate feedback parameter, for which there is evidence of genuine time dependence in GCMs [Senior & Mitchell, 2000], the TCR itself is a time-independent quantity.

For GCMs to lie along a straight line in figure 4.8 the ratio $\kappa : F_{2\times}$ would need to be constant between GCMs. For a given TCR, the RWF (or ECS) would be determined by the ratio $\kappa : F_{2\times}$ (equation 4.5.1). It may be productive for future studies to examine whether this absence of GCMs with ECS significantly below 2K is physically-based via this ratio of more statistically independent climate parameters.

## 4.6 Conclusions

This chapter has focused on interpretations of the climate response via the use of simple energy-balance representations. Whilst such models do not have the complexity of GCMs, they are much less computationally intensive than GCMs and therefore are of much greater use for policy analysis.

In this chapter I have mapped the relationship between a 2-box climate model and the commonly used metrics of climate response TCR and ECS. Using this formalism, it is possible to show that uncertainty in the TCR is the more important metric of uncertainty for GMST changes under representative future emissions scenarios and that independent uncertainty in the ECS is only of secondary importance (at long timescales or if RWF is very low). I have described a simple method of connecting uncertainty in historical forcing to uncertainty in TCR in order to produce climate projections under a range of climate responses that are consistent with observed climate change.

This chapter subsequently uses the same impulse-response energy-balance model to derive confidence intervals for the TCR (the most policy relevant metric of climate response) from observations of the 2000s energy budget. I show that the 95% confidence interval on TCR from identical observations is highly dependent on the exact model formulation (model structure) and cannot be considered 'model-independent'. Uses of the CCR approximation to estimate TCR from observations underestimate the 95% confidence interval of TCR compared to an impulse-response model calibrated to a typical GCM response. The 95% confidence intervals for the ECS estimates are strongly dependent on the assumed efficacy of ocean heat uptake, questioning whether observationally-based studies can place any meaningful constraints on the uncertainties in ECS.

Despite this sensitivity of the 95% TCR confidence interval on the response timescales in the impulse-response formulation, the 5% TCR confidence interval (and the best-estimate) seems to be robust under the same variations of the response timescales. When compared to the GCM climate response (using the more independent metrics of TCR and RWF - the ratio of TCR to ECS), observationally-based estimates are consistent with regions of low climate response that are not sampled by GCMs in either the PPE of chapter 3 or the CMIP5 MME. As shown clearly in figure 4.7, ECS and TCR are correlated in GCMs, and assuming they can vary independently does not reflect climate response in current GCMs. This offers further evidence for the value of summarising future GMST response using the more independent quantities RWF and TCR.

It must be noted that comparisons between GCM sensitivities from $CO_2$-only integrations with sensitivities derived from the historical observational period, in which multiple forcing agents have been forcing the climate system with inhomogeneous spatial coverage and different efficacies (the relative efficiency of a unit of radiative forcing from a gas in driving GMST change relative to a unit of $CO_2$ radiative forcing), are potential explanations of this low sensitivity discrepancy between the two methods [Shindell, 2014; Kummer & Dessler, 2014]. Additional caveats are the limitations of the simple climate model used to analyse the observational data. The simple model used contains no representation of internal variability in GMST, which may conflate decadal scale variability with forced response and has knock-on effects for estimates of climate response [Huber & Knutti, 2014]. A particularly important limitations of these energy-balance models is that the climate feedback parameter is assumed to be constant in time, which is seemingly contradicted by the CMIP ensemble [Winton *et al.,* 2010].

Accounting for these simplifications in the simple model, may be partly responsible for the RWF discrepancy between the observational and simple model RWF ranges and the CMIP5 ensemble. However, as no systematic attempts have been made to explore the regions of GCM climate response space consistent with a low TCR and a high RWF, it is unclear whether GCMs cannot simulate this region of climate response space consistent with observational and simple model estimates, let alone whether this inability is physically-based. The physical plausibility of a GCM with a low climate response will form the subject of the following two chapters of this thesis.

# CHAPTER 5

# Constructing a low ECS perturbed physics ensemble

The previous chapter highlighted a region of low climate response, consistent with observationally-based energy balance model estimates, that cannot be ruled out under energy-balance model structural uncertainty and not sampled by GCMs. This chapter investigates the possibility of rectifying this discrepancy between these methods of estimating uncertainties in climate response by attempting to construct low climate response GCMs. Section 5.1 first reviews postulated mechanisms for low climate response in GCMs, before a statistical emulator of the HadCM3 climate response space is developed in section 5.2. This emulator is trained on the results of chapter 3 (section 5.3), validated, and then used to identify parts of parameter space that are emulated to have low climate response and radiatively-balanced control climates (section 5.4). The chapter then goes on to describe the properties of this emulated low climate response PPE, and a sub-set of these GCMs that are simulated have ECS <2K (section 5.5).

## 5.1 Postulated mechanisms for low ECS

A variety of different mechanisms and processes can contribute to the sensitivity of the climate under an increase in atmospheric $CO_2$. Many of these mechanisms were reviewed in chapter 2, which described how the most important uncertainties in ECS and TCR are associated with cloud feedback processes (see section 2.3.3).

Chapter 3 showed that in a diverse PPE of HadCM3, there was an almost complete lack of PPE members with ECS below 2K. This result is replicated in the CMIP5 MME of GCMs from modelling centres around the world. Chapter 4 showed that there is likely a region of low TCR and high RWF that is consistent with observationally-based energy-balance estimates of the climate response, based on observations of the 2000s, and which is not sampled by GCMs. Independent of these observationally-based energy-balance estimates of climate response, this seemingly 'hard' lower limit for ECS in GCM ensembles requires an investigation into its physicality. Is this limit a physical limit, telling us it is implausible to construct a GCM with ECS <2K that resembles the real world? Or is it merely a consequence of insufficiently diverse GCM ensembles which do not sample the plausible climate response space fully (see Tebaldi & Knutti [2007] for a discussion of the challenges in interpreting multi-model ensembles)?

The climate response, both TCR and ECS, are emergent properties of a GCM, meaning they have complex dependencies on the structure of the GCM and the values of the sub-gridscale parameters (see chapter 3) and hence cannot easily be controllably perturbed [Mauritsen *et al.,* 2012]. Additionally, ECS is not a variable that GCMs are explicitly evaluated against in the GCM development cycle, largely because it is not a directly observable property in the current climate system. GCMs are however evaluated against the observed data from the historical period and, as shown in figure 4.1, varying assumptions about historical aerosol forcing show that a range of different future climate responses could be consistent with the historical record of GMST, planetary heat uptake and present-day radiative forcing uncertainty.

Despite the challenge in explicitly tuning GCMs to simulate a specified ECS, there are postulated mechanisms that could lead to very low ECS and recent attempts to simulate them in a GCM. The most well-known of these mechanisms is the so-called "Iris" hypothesis [Lindzen *et al.,* 2001]. This mechanism postulates a decrease in cirrus cloud detrainment from tropical cumulonimbus clouds as surface temperatures increase. The reduced spatial coverage of these clouds, and hence a reduction in the LW-CRE of these high clouds, coupled with a drying of the free-troposphere beneath them, would act as a negative feedback as the climate warms under the direct $CO_2$ greenhouse effect. Lindzen *et al.* [2001] claim evidence for this mechanism in observations of the "warm-pool" region of the western Pacific Ocean boundary, but this conclusion has been disputed by other work (e.g. Hartmann & Michelsen [2002]). Recently, Mauritsen & Stevens [2015] explicitly parametrised an Iris effect with the ECHAM6 GCM by introducing a dependency on surface temperature into the conversion of cloud water to rain, in an effort to promote an increase in the efficiency of precipitation within convective towers and hence increased drying of the adjacent subsidence regions. This parametrisation has a modest effect on the ECS, reducing the estimated ECS by approximately 0.4K, due to compensating increases in SW cloud feedbacks associated with a general reduction of cloud fraction (due to drying of the atmosphere) all across the globe. However, Mauritsen & Stevens [2015] hypothesise that if an Iris effect could be generated by a GCM without a compensating increase in global-mean SW cloud feedback, the most uncertain component of cloud feedback, ECS as low as 1.2K could be achieved.

This chapter aims to investigate the possibility of simulating low ECS in a GCM, without introducing new mechanisms that do not arise from within the current model structure. Instead of explicitly introducing and testing hypothesised physical mechanisms, either an Iris effect or an alternative, I instead seek to emulate regions of the HadCM3 parameter space that are consistent with low ECS and plausible control climates. By exploring dependencies of ECS on sub-gridscale parameters and systematically constructing a mapping of these dependencies, a broader sense of the possible diversity and limits on plausible ECS within the GCM structure can be gained, without directly requiring artificially-introduced and therefore possibly non-physical mechanisms.

## 5.2 Constructing a simple multivariate emulator for climate response

An emulator of a GCM is a predictor of the GCM properties as a function of the GCM sub-gridscale parameters [Rougier, 2007]. This can be represented as a mapping between the vector of independent sub-gridscale GCM parameters, $\mathbf{x}$, and some vector of GCM simulation statistics, $\mathbf{y}$,

$$\mathbf{y} = f(\mathbf{x}) + \epsilon, \tag{5.2.1}$$

where $\epsilon$ is the error term associated with the mapping, $f$. $f$ can be evaluated for specific points in parameter space by conducting a climate simulation with a specific input parameter set. However, as finite computing resources limit the number of feasible simulations of points in parameter space, a statistical model trained on the simulated points in parameter space is required to make $f$ continuous across the multi-dimensional GCM parameter space (assuming that the GCM properties are a reasonably smooth function of changes in the input parameters).

Emulation of climate response from a PPE has been attempted in several previous studies (e.g. Murphy *et al.* [2004] and Sexton *et al.* [2012]). However, unlike in these studies, which focused on quantifying the distribution of ECS as their stated aim, the emulator I construct here will instead be used to predict a low climate response ensemble. Targeting uniquely low climate responses within the HadCM3

GCM structure will likely involve extrapolated predictions for the climate response at parameter values outside of the ranges of the ensemble described in chapter 3 (which will be used to train the emulator). Therefore, the statistical emulator will be used only as an approximate guide to interesting points in parameter space and explicit simulation of the climate response at these locations will be required to confirm the skill of the statistical emulator. I construct an emulator based on the structure described in Sexton *et al.* [2012]. This emulator uses multivariate linear-regression to provide a direct mapping between the input parameters and several statistics of the GCM climate, both in the control state and under increased concentrations of atmospheric $CO_2$. The expected value for a scalar statistic of a GCM simulation, $Y$, is emulated as,

$$\mathcal{B}(Y|\lambda, \theta) = \sum_{k=1}^{p} \gamma_k g_k(l(X)) + \nu, \tag{5.2.2}$$

where $p$ is the number of independent regressors, $\gamma_k$ are the regression coefficients, $g_k$ are the regressors (functions of $l(X)$, themselves transformations of the underlying GCM parameters, $X$) and $\nu$ represents Gaussian random noise errors. $\mathcal{B}$ represents a Box-Cox power transformation of $Y$, indexed by the parameters $\lambda$ and $\theta$. Key features of the emulator are described in the subsections below.

### 5.2.1 Transformations of regressors and transformations of dependent variables

Polynomials of the untransformed GCM parameters may not be the best functional form with which to predict the dependent variables of interest. In particular, ECS has been shown to be better predicted by logarithmic transformations of the GCM parameters in a previous emulator [Sexton *et al.*, 2012]. Whilst, for discrete parameters I restricted regressors to be the untransformed GCM parameters, $l(X) = X$, I allow continuous parameters to be either untransformed ($l(X) = X$) for all continuous parameters or logarithmically transformed ($l(X) = \ln(X)$) for all continuous parameters, creating two different potential sets of regressor transformations for the emulator of a particular dependent variable.

The emulator structure assumes a Gaussian distribution for the regression model error, and most common methods of statistical inference typically rely on the validity of this assumption. If this assumption is strongly violated with the data and statistical model considered, significance tests can give potentially misleading results and under-estimates of statistical uncertainties. In equation 5.2.2, $\mathcal{B}$ represents a Box-Cox transformation (see Draper & Smith [2014] for a full discussion), a common power transformation used to ensure that the residuals of the regression are as close to Gaussian as possible. The Box-Cox transformation is given as a function of the parameters $\lambda$ and $\theta$ by,

$$\mathcal{B}(Y|\lambda,\theta) = \begin{cases} (sy_i - \theta)^{\lambda} & \text{if} \quad \lambda \neq 0 \\ \ln(sy_i - \theta) & \text{if} \quad \lambda = 0, \end{cases} \tag{5.2.3}$$

where $\lambda$ is the Box-Cox parameter, $y_i$ is a realisation of the GCM climate variable $Y$, $s$ is allowed to take values of $\pm 1$ and $\theta$ is a parameter chosen to ensure that $(sy_i - \theta) > 0$ for all $i$, a required condition of a monotonic transformation.

For each transformation of continuous regressors ($l(X) = X$ or $l(X) = X$) and value of $s$, a set of regressors is selected using the automated step-wise method described in section 5.2.2. For this regression model, a value of $\lambda$ (rounded to the nearest half integer) is chosen to maximise the log-likelihood of the dependent variable given the parameters of the regression model (see figure 5.1 for an example with the ECS emulator). I select a single best combination of regressor transformation and value for $s$ by selecting the combination with the minimal residual sum of squares between the fitted values for $Y$ and the training data in the original units of $Y$.

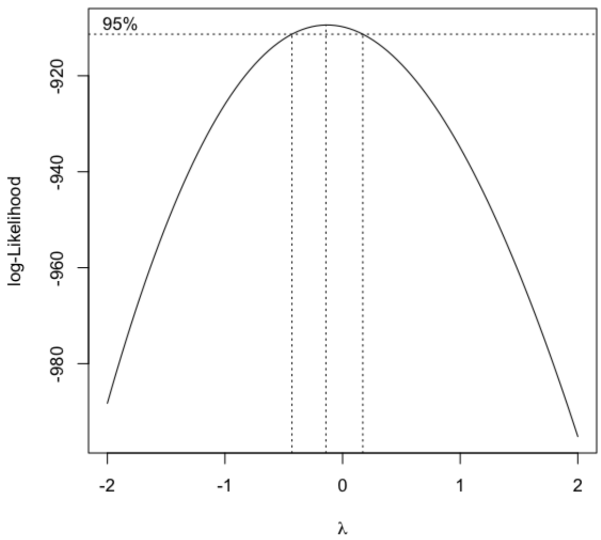


**Figure 5.1** – Box-Cox plot of regression model log-likelihood for the fit between the step-wise selected regression model (using logarithmic transforms of GCM parameters) and the data for different values of the parameter $\lambda$. $\lambda = 0$ corresponds to a logarithmic transformation of the dependent variable.

### 5.2.2 A step-wise automated regressor selection method

To compare statistical models with different regressors a common measure of model predictive skill is required. The Akaike Information Criterion (AIC) is a commonly-used metric of model skill, with the advantage that increases in predictive skill (increases in log-likelihood) are weighted against increases in regression model complexity (which helps to guard against overfitting) via a penalty term proportional to the number of independent parameters. The AIC is an approximate measure of the *relative* loss of information (compared to knowing the entire underlying data set) by selecting one regression model as opposed to another [Akaike, 1974]. Mathematically, the AIC for a regression model is defined as,

$$\mathrm{AIC} = 2k - 2\ln L, \tag{5.2.4}$$

where $k$ is the number of independent regression model parameters and $L$ is the maximum of the regression model's likelihood function (the likelihood of observing the data given the best values of the regression model's parameters). A regression model with a lower AIC contains more information than a regression model with greater AIC.

As 32 independent parameters are perturbed in climate*prediction*.net PPEs, a very large number of possible regressors exist when allowing for the possibility of important interaction effects (the effect of variations in one parameter dependent on the value of other parameters) between parameters. This large number of regressors necessitates the use of an automated model selection algorithm in order to remove statistically redundant terms and reduce the risk of overfitting. I use an automated both ways step-wise algorithm, in which (from a maximal set of candidate regressors) all possible regressors (within the scope of terms allowed) are separately added or removed to/from the current regression model and the AIC assessed for the maximum-likelihood fit of each resulting regression model. The regression model with the minimal AIC out of all the candidate regression models is then selected from this step and the process of adding/removing all possible terms is repeated from this new working regression model. This process is continued until the addition/removal of any terms will result in a regression model with a higher AIC than the working regression model. Whilst automated regressor selection algorithms have their drawbacks, particularly the inability to inspect and test the working regression model at each stage of the process, they form the only practical method of deriving a "best" set of regressors for the emulator under such high dimensionality.

As a maximal set of regressors I consider two-way interactions between transformations of all GCM parameters as well as regressors that are quadratic and linear in the transformed GCM parameters ($l(X)^2/l(X)$). Only a subset of the possible three-way interactions are allowed in order to prevent the selection algorithm becoming too computationally intensive whilst also allowing the emulator scope to include three-way interactions between cloud and convection parameters that have been previously shown to have a significant role in predicting ECS, namely *entcoef*, *vf1*, *ct*, *cw-land*, *rhcrit* and *eacf* [Rougier *et al.*, 2009]. For computational simplicity, GCM parameters that are not continuous (*start_level_gwdrag*, *r_layers* and *diff_exp*) are only considered as linear terms without any potential interactions.

## 5.3 Fitting and validating the emulator

The emulator is fit, as described in section 5.2, to the PPE of chapter 3. GCM properties emulated are ECS, TCR, control climate global-mean reflected shortwave radiation (RSR) and control climate global-mean outgoing longwave radiation (OLR). Averages over the final 20 years of the control integration are used as training data for the emulators of RSR and OLR. The emulator selects logarithmic transformations of the continuous GCM parameters for ECS and linear transformations for TCR, RSR and OLR. $\lambda = 0$ (logarithmic) is selected as the Box-Cox transformation parameter for ECS, consistent with other emulator studies that have investigated ECS [Rougier *et al.*, 2009; Sexton *et al.*, 2012]. TCR takes a transformation of $\lambda = 0.5$, indicating a raw distribution similar to, but not quite, a log normal distribution. RSR requires $\lambda = -2$ and OLR requires $\lambda = 2$, indicating opposing senses of skew in the raw distributions, as might be expected due to the imposed condition of approximate radiative balance on the training data. Regressors are transformed with logarithmic, linear, logarithmic and logarithmic transformations for ECS, TCR, RSR and OLR respectively. $s = 1$ for all four emulators.

A number of different tests are conducted in this section to validate the emulators. The regression model shown in equation 5.2.2 assumes that the distribution of residual errors after fitting can be represented by a Gaussian distribution of constant variance at all points in parameter space. The validity of such an assumption can be investigated by a comparison of the residual error distribution against a Gaussian distribution, commonly known as a quantile-quantile (q-q) plot, as shown in the panels of figure 5.2. The y-axes shows the quantiles (each member of an $n$ member

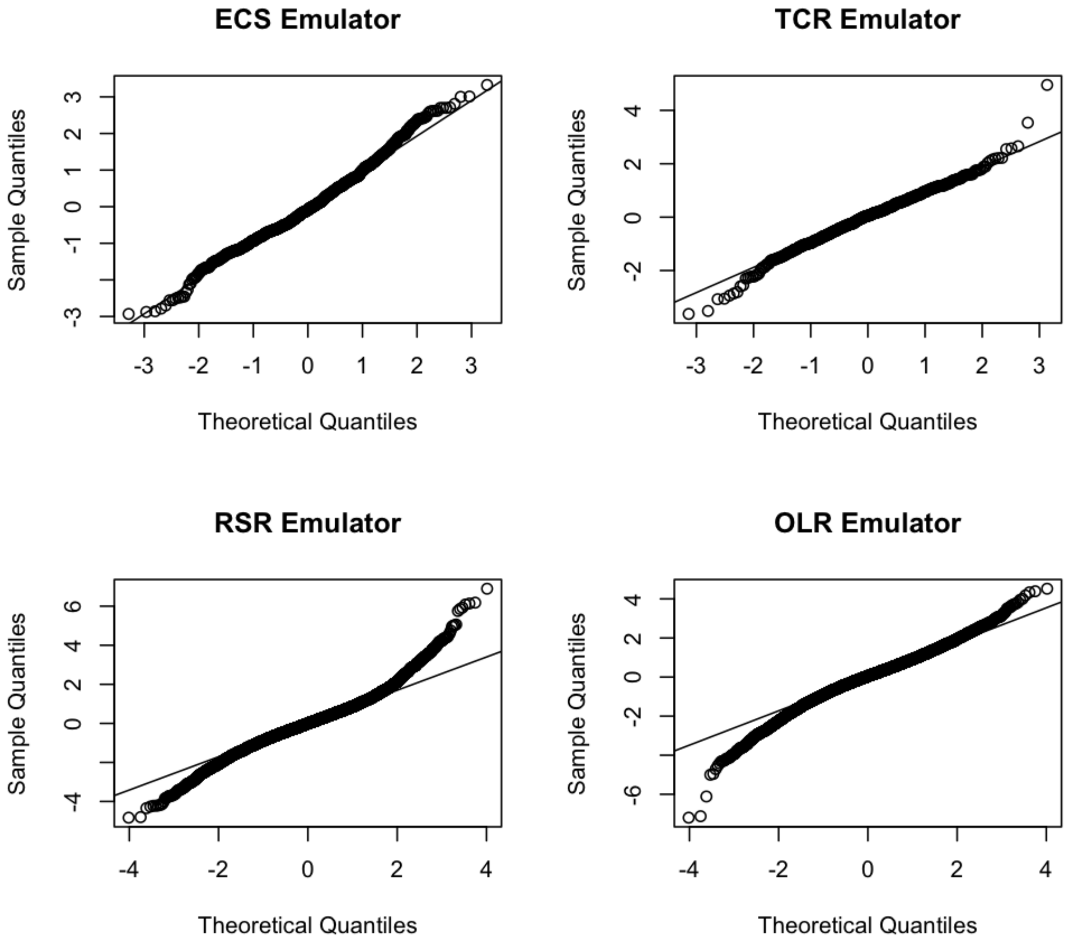


**Figure 5.2** – Quantile-quantile (q-q) plots. The panels show the emulators for ECS, TCR, RSR and OLR respectively. The y-axis shows standardised errors of the emulated predictions, centred and normalised using an estimated standard deviation, against the quantiles of a Gaussian distribution with mean of zero and unit variance (x-axis). The thin balck line shows the 1:1 line.

distribution represents one of the $n$ equally likely quantiles of the distribution) of the residuals of the emulated fit normalised by their standard deviation. The x-axes represent the location of the corresponding quantile in a normal distribution with unit standard deviation and zero mean. Residual distributions consistently close to the 1:1 line in the panels of figure 5.2 indicate strong Gaussianity across the distribution. The step-wise selected linear model for ECS and TCR generate a generally good approximation of Gaussian errors across the distribution. For RSR, there is a substantial deviation from Gaussianity in the upper tail of the distribution, with a greater number of >4 standard deviation underestimates of the data than would be expected from a Gaussian distribution of residuals. A corresponding anomalously large number of >4 standard deviation overestimates (which may be expected due to the constraints of approximate radiative balance that is imposed on the training data), is observed for the OLR emulator. These deviations from Gaussianity are indicative of the opposite sense of skew (due to energetic imbalance constraints) in the training data of these two distributions (that clearly cannot be totally corrected by a Box-Cox transformation). As these deviations from Gaussianity fall well outside the 5-95 percentiles (approximately equivalent to ±2 standard deviations), they are unlikely to have any substantial impact in the prediction of 5-95 percentile confidence intervals on predictions of low climate response GCMs near the centre of the RSR and OLR distributions.

Quantile-quantile plots are not in-themselves indicative of emulator predictive skill and only serve to validate the regression model structure and statistical assumptions. Leave-one-out (LOO) diagnostics [Rougier *et al.,* 2009] indicate the scale of predictive uncertainty arising from the emulator when a single data point of the training data is removed and predicted using all other data points. If the emulator corresponds to a well-fit regression model, the actual value for the predicted data point should lie outside of the 95% prediction intervals (which account for uncertainty in the regression coefficients as well as the contribution to the prediction uncertainty associated with the statistical Gaussian errors) roughly 5% of the time.

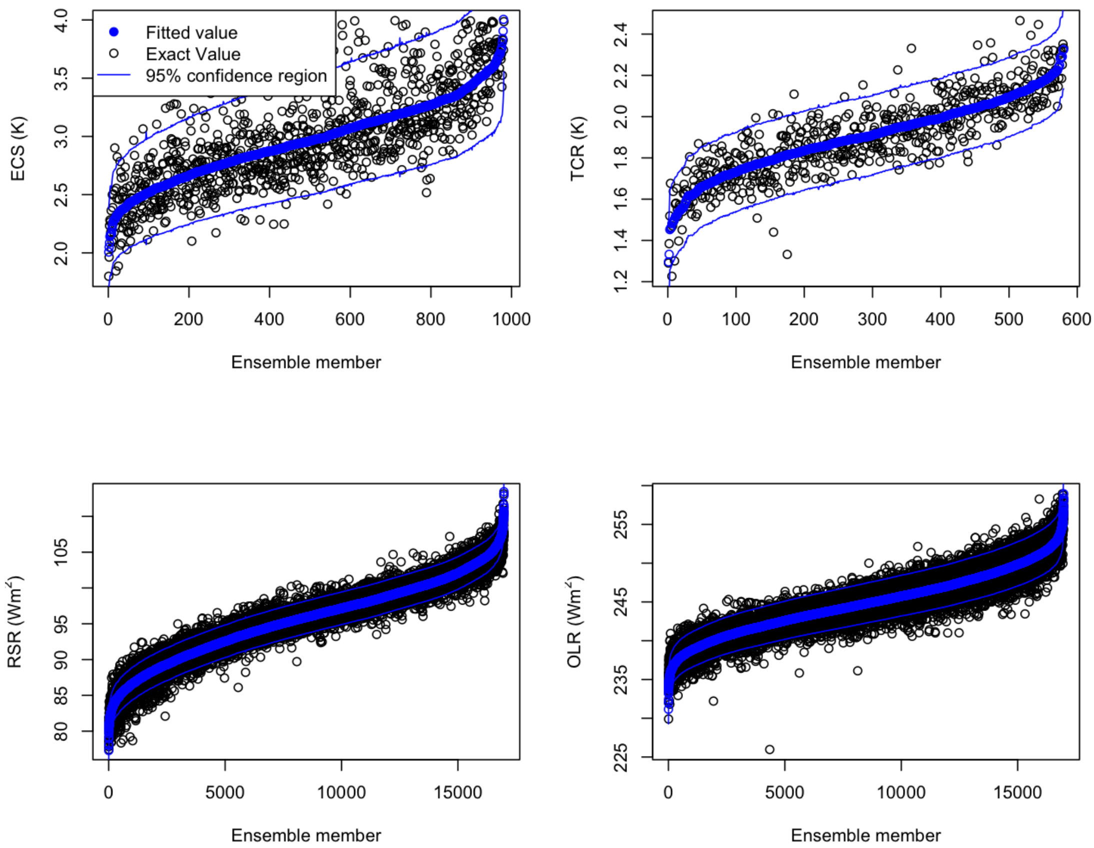


**Figure 5.3** – The panels show LOO cross-validation for the emulators of ECS, TCR, RSR and OLR respectively. Blue filled circles show the ordered emulated values for each member of the training ensemble, when the emulator is trained using all ensemble members except the one in question. Simulated values (from chapter 3) are shown in each case by displaced black circles at the same x-axis position. Blue lines show the 95% prediction intervals for each fitted value.

If more (less) than 5% lie outside of the prediction intervals, predictive uncertainty is under(over)-estimated by the emulator.

Figure 5.3 shows LOO validation for the emulators of ECS, TCR, RSR and OLR. In all cases roughly 5% of ensemble members lie outside of the 95% prediction intervals, with no statistically significant differences from the 5% fraction for any of the four emulators (e.g. 5.6% lie outside in the case of ECS emulator). The 95% prediction intervals given by the LOO exercise are likely to be a good indicator of the predictive uncertainty from the full emulator as one data point is a small fraction of the total ensemble size. LOO uncertainty intervals are noticeably larger in the

ECS emulator than the TCR emulator, despite the advantage of a larger ensemble size (more experiments completed the STEP as opposed to RAMP experiment for the training ensemble), partly reflecting the non-negligible uncertainty inherent in the regression-based method with which the ECS values for the training ensemble have been calculated. 5-95% predictive uncertainties are typically on the order of 1K for the ECS, emphasising the need for very large ensembles in order to produce reasonably sized ensembles that actually do satisfy intended constraints on ECS.

A stronger test of the emulator than a LOO test is a traditional cross-validation exercise. A leave-out-one-third (LOOT) test removes a random one third of the data points, fits the emulator on the remaining two thirds of the data, and then predicts the value of the missing data points. If the emulator is well-fit the studentized residuals (normalised by an independent estimator of the standard deviation of the regression model errors) at the predicted data points should form a Student-t distribution with the residual degrees of freedom.

As shown in figure 5.4, whilst in all cases there is some mis-fitting, all emulators display LOOT residual distributions that are good approximates of the appropriate Student-t distributions. The evidence from this section, and in particular the good performance of the emulators under this more stringent test of their predictive power, offer support to the statistical assumptions and methodology that underlie the statistical emulators of GCM response. The emulators provide a good guide to the variation throughout the training ensemble and are useful tools to search for parts of parameter space that correspond to realistic climates with low climate responses.

Examining the impact of variations in important GCM parameters on emulated ECS can offer physical insights into possible sources of ECS uncertainty. Figure 5.5 shows the effect of variation in all continuous parameters included in the final set of regressors for the ECS emulator when a particular parameter is varied across its range in the training ensemble whilst holding all other parameters fixed at their

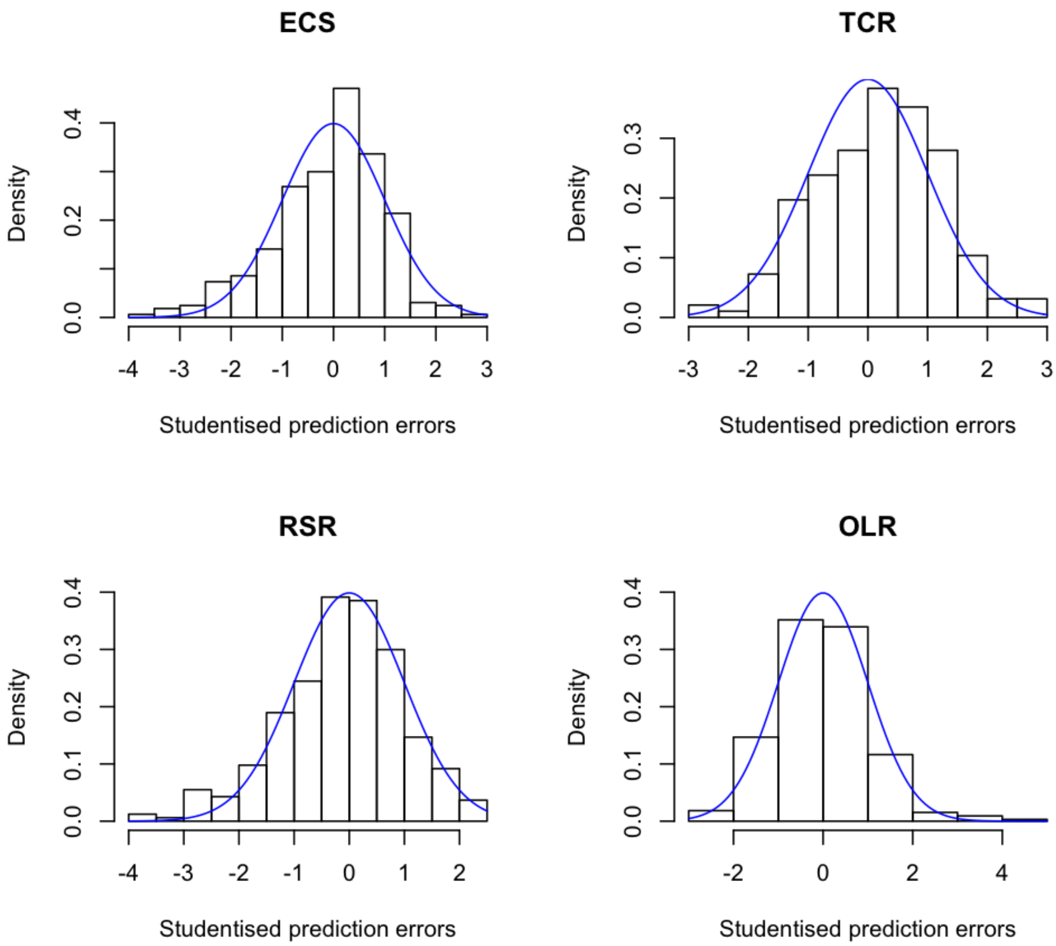


**Figure 5.4** – Leave-out-one-third (LOOT) cross validation. Each panel shows a histogram (in black) of studentised prediction errors for a random one-third of the members of the training ensemble using the emulator trained on the other two-thirds. The blue curve shows a Student-t distribution with the residual degrees of freedom.

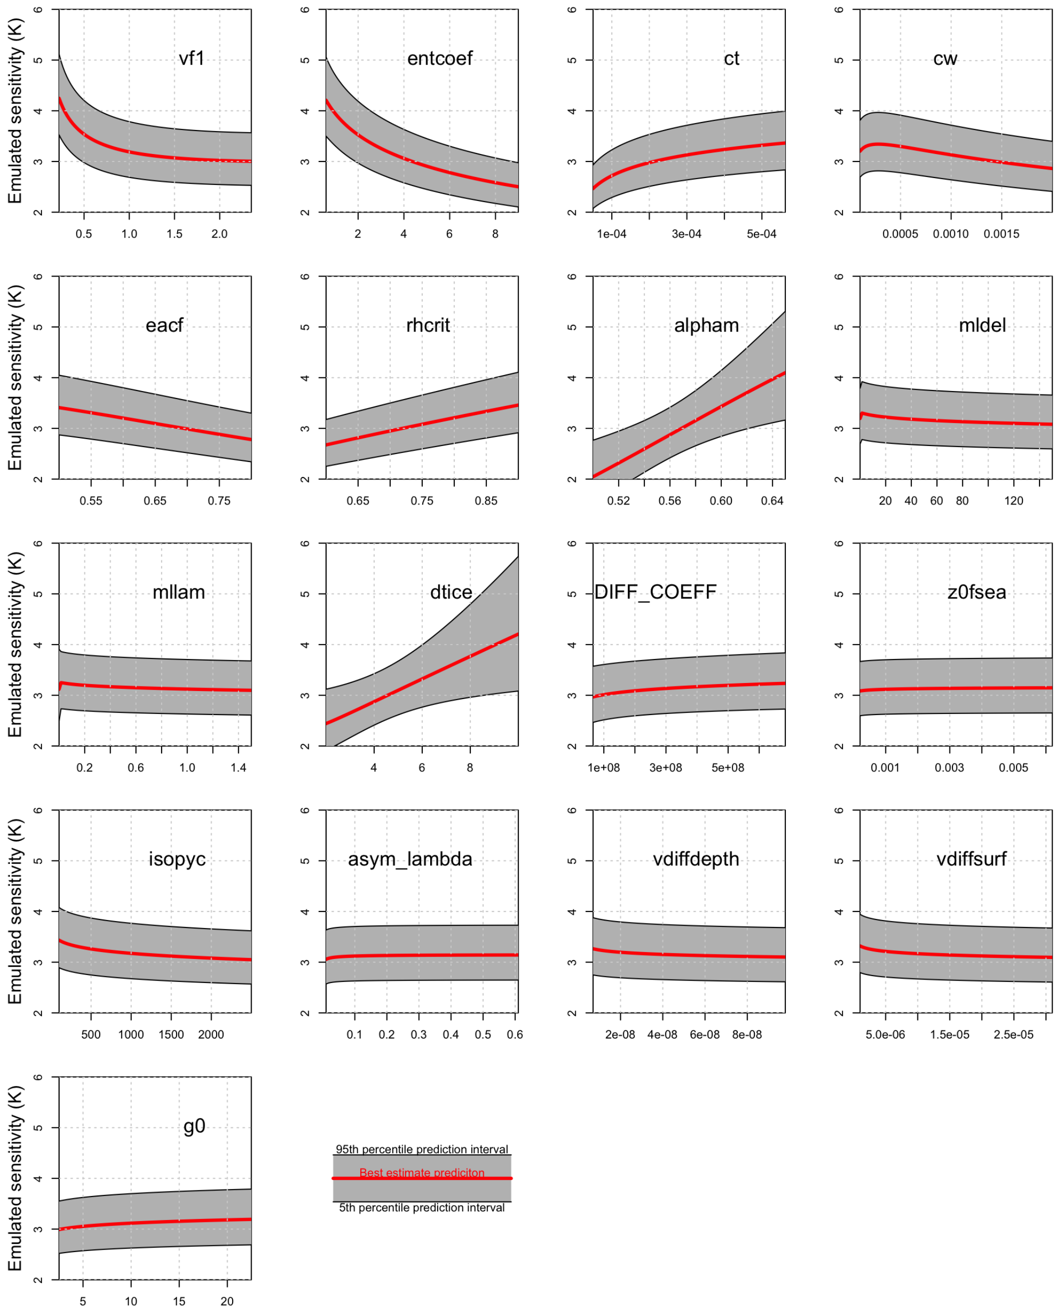


**Figure 5.5** – Effect of variation in continuous parameters on emulated ECS. In each panel the variable considered is varied across the entire range in the training ensemble whilst all other parameters are fixed at their ensemble mean value. The red line shows the best-estimate prediction of the ECS whilst the dark shaded bands show the 5-95% prediction intervals.

ensemble mean values. Interaction effects between parameters in the emulator are not captured in figure 5.5 but only act to scale the magnitude of the effects shown and do not affect the sign or shape of the effects. The emulated effects of the parameters on ECS within this ensemble are in general consistent with a previous study [Rougier *et al.,* 2009] that looked at the role of sub-gridscale parameter variation on ECS within the HadSM3 GCM (a slab ocean configuration of the HadCM3 GCM, sharing the same atmospheric model). Despite the large prediction errors associated with statistical and internal variability uncertainty, the combined interactions of statistically significant sub-gridscale parameters are sufficient to clearly distinguish certain regions of parameter space from one another in terms of their emulated climate response (not shown). In particular the emulator highlights low values of *entcoef* and *vf1* as driving very high climate sensitivities in the ensemble, consistent with other studies of PPEs [Murphy *et al.,* 2004; Sanderson *et al.,* 2008]. The impact of the entrainment coefficient (*entcoef*) has long been known to drive high climate sensitivities within HadSM3 ensembles, due to the positive LW feedback of the enhanced stratospheric water vapour content that results when ambient mixing of a convecting parcel is minimal in areas of deep convection [Sanderson *et al.,* 2008]. Parameters that control the reflective properties of sea-ice, *dtice* (the temperature range that sea-ice albedo varies over) and *alpham* (the sea-ice albedo at point of melting) are emulated to have impacts on ECS. Increasing the sea-ice albedo at melting creates a strong amplifying SW clear-sky feedback as a larger albedo decrease is experienced when the sea-ice melts under $CO_2$-induced warming, assuming all else is equal. Interestingly, no significant correlations is observed between clear-sky SW feedbacks and *alpham* or *dtice* in the training ensemble, possibly suggesting significant co-dependencies with other parameters in order to satisfy the imposed global RSR constraints.

Variations in *ct*, *cw-land* and *rhcrit* are all parameters that could be important to the simulation of control climate cloud cover and cloud feedbacks. Increasing *ct* acts to reduce the rate at which liquid cloud water content increases (increased $c_T$ in equation 3.1.3), which contributes to less lower cloud cover. Similarly, increases in *rhcrit* increases the gridbox relative humidity required for cloud formation, reducing cloud cover all other things being equal. Control climate cloud properties have been shown to be important for the climate response to a $CO_2$ perturbation (e.g Webb *et al.* [2006] or Yokohata *et al.* [2010]), indicating that similar relationships are likely to be apparent within this PPE.

When interpreting the emulated impact of variations in a parameter on the ECS in this PPE, it must be remembered that the training ensemble used to fit this emulator displays some colinearity of parameters, due to the constraints on RSR and OLR imposed on the training ensemble of chapter 3. Colinearity can make the partition of regression coefficients between colinear terms sensitive to the exact formulation of regression model. Therefore, it is important not to over-interpret the individual regression coefficients within the emulator, but to focus on the sign and shape (and potential underlying physical mechanisms) of the effects, when interpreting the emulated dependencies of ECS on the parameter values.

## 5.4 Predicting a low sensitivity ensemble

The emulator developed in the previous section can now be used to identify parts of GCM parameter space that have a low climate response and radiatively-balanced control climates. The two following subsections briefly describe the search criteria for climate response and control climate simulation properties that are targeted with the emulator.

### 5.4.1 Target for low climate response

I use the ECS emulator to search for points in GCM parameter space that have ECS<2K. I select for low climate response using the ECS emulation only (as opposed to additionally using TCR/RWF) in order not to impose too many constraints on the linear emulator. As TCR is better constrained by observations (as shown in chapter 4) ECS is therefore a more interesting target for GCM experiments. By emulating ECS, as opposed to any of its constituent forcing or feedback components, no specific constraints on the mechanisms by which the GCM may achieve a low climate response are *a priori* imposed. Low ECS could, in theory, be achieved through negative (or weakly positive) climate feedback in any of the components of the climate feedback parameter or via substantial fast tropospheric adjustments that result in a low effective radiative forcing of $CO_2$.

### 5.4.2 Targets for GCM control climate radiative fluxes

An essential requirement for emulating plausible climate simulations in PPEs is a realistic pre-industrial control climate. For low ECS GCMs to be indicative of realisable climate response in the real climate system, constraints on some aspects of the GCM climatology need to be imposed. Many possible constraints could be placed on GCM simulation to capture some aspect of fidelity with the real world (e.g. Murphy *et al.* [2004] and Williamson *et al.* [2013]). As I am attempting to explore the limits of possible climate response uncertainty, I choose to place only simple constraints, via an estimate of observational-model discrepancy, on the TOA radiative fluxes (RSR and OLR) in the GCM control climate. GCM fidelity can be assessed through any number of comparisons with observations and reanalysis, however as the PPE members predicted by the emulator will not have been tuned to improve their climate simulation (without significantly altering their ECS), but in theory could be; applying more conditions to the emulation of the control climate may lead to an overly restrictive search of parameter space.

To assess pre-industrial radiative fluxes of PPE members an estimate of "observational-model discrepancy" is required [Tett *et al.,* 2013]. Such discrepancy includes uncertainty in the observations (described by a covariance matrix $\sigma_{OBS}$) and intrinsic uncertainties associated with the modelling of the observational period (described by a covariance matrix $\sigma_{MOD}$). I follow Tett *et al.* [2013] and Yamazaki *et al.* [2013] to estimate the total uncertainty covariance, a brief summary of the method is provided here, but readers should consult Tett *et al.* [2013] for full details. An observational uncertainty covariance is estimated for the joint distribution of RSR and OLR using the CERES satellite dataset [Loeb *et al.,* 2009] over the period 2000-2005, given as,

$$\sigma_{OBS} = \begin{pmatrix} 0.74 & -0.64 \\ -0.64 & 0.82 \end{pmatrix}. \tag{5.4.1}$$

The existence of off-diagonal terms indicate the covariance between RSR and OLR in the observational timeseries when the energetic imbalance of the planet over the observed period is taken into account, as an increase in either RSR or OLR is expected to be compensated by a decrease in the other in order to preserve approximate radiative balance. These terms are not perfectly anti-correlated, representing the small variations in the incoming solar radiation over the 2000-2005 period as well as the net energetic imbalance of the climate system (estimated to be 0.75Wm$^{-2}$ by Tett *et al.* [2013]).

Additional modelling uncertainties are assumed to be Gaussian and independent for simplicity. These uncertainties aim to capture additional factors that may cause discrepancies between GCMs and observations over the observational period, but do not necessarily indicate an incompatibility of the GCM (and therefore parameter set) with the observations. Factors included here are initial condition uncertainty and uncertainty in the correct radiative forcing of the climate system for the 2000-2005 period (again see Tett *et al.* [2013] and Yamazaki *et al.* [2013] for full details). The estimated total covariance matrix, $\sigma_{TOT}$ ($\sigma_{TOT} = \sigma_{OBS} + \sigma_{MOD}$), is given as,

$$\sigma_{TOT} = \begin{pmatrix} 7.6 & -4.5 \\ -4.5 & 4.3 \end{pmatrix}. \tag{5.4.2}$$

Many subjective choices are involved in estimating this total covariance matrix for assessing "observational-model discrepancy", including the assumption made here that this uncertainty covariance matrix, estimated for the 2000-2005 period, also accurately describes total observational-model discrepancies for the pre-industrial climate. Without observations of the RSR and OLR in the pre-industrial climate, I choose to centre this uncertainty covariance around the standard HadCM3, due to its close approximation to a radiatively-balanced control climate. There are additional sources of "observational-model discrepancy", such as radiative-transfer model uncertainty, that have not been included here, and therefore the precise values in the total covariance matrix should not be over-interpreted, but can act as an approximate guide to the real-world consistency of radiative fluxes in the GCM control climates. Approximate radiative balance, and therefore a non-drifting climatological GMST, is no guarantee of a good representation of the climate in all aspects. However, selecting for emulated low ECS points in parameter space that lie within the 99% total-uncertainty confidence interval around the standard HadCM3 allows diversity of control climate and climate response to be maintained whilst also restricting the search to GCMs that preserve important constraints of energetic balance in the quasi-equilibrium control climate.

### 5.4.3 Sampling methodology

Latin hypercube sampling (a space-filling method in which multi-dimensional parameter space is systematically sampled in order to fully explore the allowable range of each GCM parameter [McKay *et al.,* 2000]) can be used to search for points in parameter space that satisfy the targets described in sections 5.4.1 and 5.4.2. GCM parameter ranges from the training ensemble range are expanded by 25% at both the upper and lower ends, for all parameters except where physically bounded ranges exist (such as parameters that are positive definite), in order to maximise the number and diversity of points in parameter space that are emulated to satisfy all constraints. Large ensemble sizes are required to offset potential predictive failures of the emulator and the non-negligible statistical and internal variability uncertainty in the emulated ECS. From a 100,000 member sampling of these parameter ranges 6400 parameter sets are emulated to satisfy both climate response and radiative constraints. A random 500 of these were simulated in what I will refer to as the predicted low sensitivity ensemble.

Unlike for the atmosphere, oceanic adjustment timescales are on the order of several centuries. Therefore perturbations to the oceanic parameters could create energetic imbalances with the available oceanic initial states that could take centuries to re-equilibriate. To avoid this, I restrict oceanic parameter variation within the sampling structure to configurations that have already been equilibriated, whilst maintaining the target conditions of the emulated response. As atmospheric perturbations above these oceans may lead to disequilibrium between the atmosphere and upper oceans, I conduct 10 years of further control integration before the STEP forcing is implemented to allow approximate equilibrium between the atmosphere and upper ocean to be re-established.

With many degrees of freedom within the set of emulators considered here, there are potentially many ways of satisfying the target constraints. In particular, the number of "effective levers" deployed by the emulator to satisfy the goal of ECS<2K, whilst additionally satisfying the constraints on the radiative fluxes in the control climate, may be informative about potential diversity in other features of the simulated climate. Diversity of emulated climate response can be estimated by calculating a total path-length in emulated ECS space between any two members of this predicted low ECS ensemble, normalised by the net difference in emulated ECS between the two members. Mathematically, this "path-length" between two members of the ensemble, $\mathcal{L}_{ij}$, can be expressed as,

$$\mathcal{L}_{ij} = \sqrt{\frac{\sum_{k=1}^{P}(\gamma_k g_k l(X_i) - \gamma_k g_k l(X_j))^2}{P}}. \tag{5.4.3}$$

Figure 5.6 shows a histogram of the relative path lengths ($\frac{\mathcal{L}_{ij}}{ECS_i - ECSj}$) calculated over all possible pairs of ensemble members. The distribution of relative path lengths is peaked around 20 (a value of 1 would correspond to all differences in parameter values between two ensemble members having the same sign of impact on the emulated ECS difference), indicating substantial cancellation between changes in ECS due to differences in individual parameters. Many pairs of ensemble members have relative path lengths well in excess of 100, indicating highly diverse parameter variations give similar net ECS under the emulator targets. High relative path lengths in ECS emulation is not a guarantee of meaningful diversity of climate response mechanisms (for instance, diversity in spatial patterns of climate change may be driven by parameter variation that isn't emulated to be important for determining ECS, a global mean quantity). However, it is likely that the emulator, if correctly distinguishing "noise" parameters from non-linearities, is not producing emulated low climate response by limiting parametric variation to one or two "effective" GCMs driven by differences in a few important parameters combined with large variations in "noise" parameters.

## 5.5 An ensemble-based search for low climate response HadCM3 PPE members

The emulator described previously in this chapter, predicts parameter sets with emulated ECS<2K, whilst maintaining approximate radiatively-balanced control climates. Whilst the emulators perform well under the validation tests considered previously, this is no guarantee of correct predictions, particularly when the predictions are extrapolations from the ranges of parameters within the training ensemble. This section describes features of the predicted low sensitivity ensemble for STEP, RAMP and pre-industrial control simulations carried out at the predicted locations in HadCM3 parameter space.

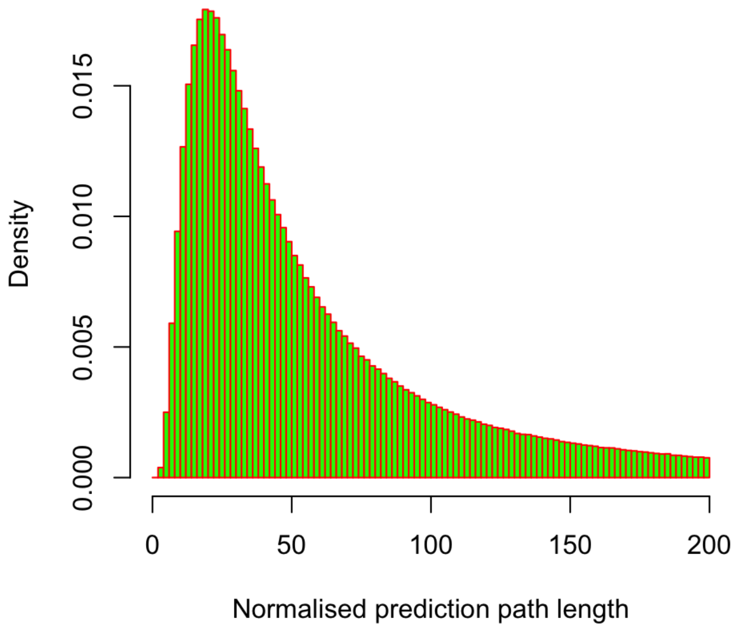


**Figure 5.6** – Histogram of relative ECS path lengths ($\frac{\mathcal{L}_{ij}}{ECS_i - ECSj}$) between all pairs of individual points in parameter space that are predicted to have ECS<2K.

### 5.5.1 Radiative biases in simulated control climates

The emulation of the predicted low sensitivity ensemble focused solely on global-mean quantities in the pre-industrial control climate. However, whilst global-mean values could be similarly simulated across the ensemble, this may hide substantial differences at regional scales. Figure 5.7 shows zonally-averaged clear-sky and cloud radiative fields for control simulations from the predicted low sensitivity ensemble, expressed as anomalies relative to CERES climatology. In general, there are few systematic clear-sky biases from CERES (except around the southern hemisphere ice-edge in SW clear-sky radiation) relative to the standard HadCM3. The correction of the standard HadCM3 SW clear-sky bias in the Arctic apparent in some ensemble members, is caused by a partial alleviation of insufficient summer sea-ice coverage in the standard HadCM3 (discussed in chapter 3 and further in chapter 6).

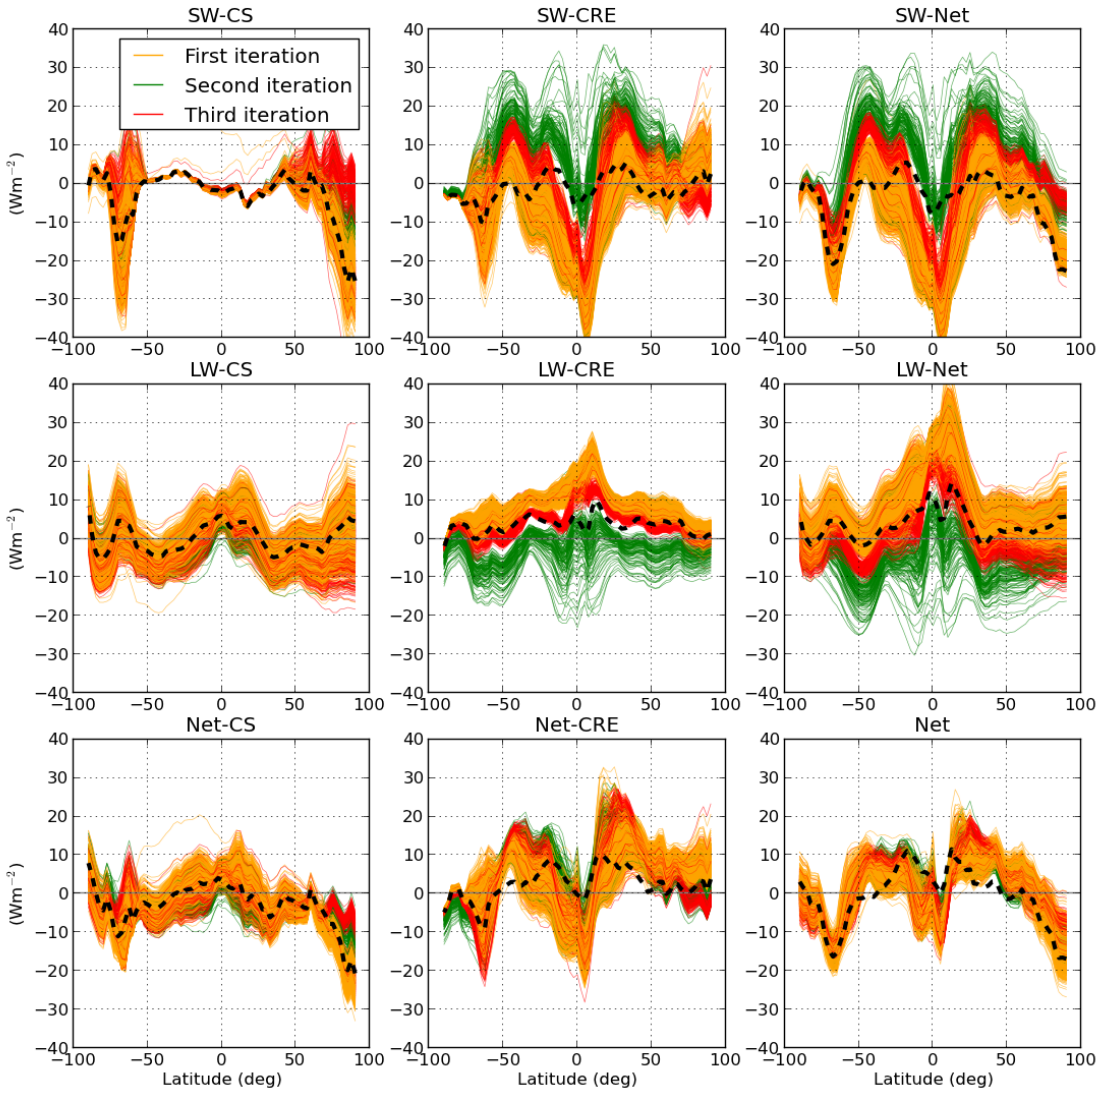


**Figure 5.7** – Zonally-averaged radiative biases in the predicted low sensitivity ensemble, shown as anomalies relative to the CERES climatology. For all fields a downward radiative flux is defined as positive. The dashed black line represents the standard HadCM3. Different iterations of the low sensitivity ensemble are denoted by colours, orange (first iteration: only global-mean radiative constraints), green (second iteration: 30°N-30°S and global-mean constraints) and red (third iteration; 60°N-40°N, 30°N-30°S, 40°S-60°S and global-mean constraints). The left hand column shows clear-sky radiances (CS), the middle shows cloud radiative effects (CREs) and the right shows combined all-sky fluxes (Net).

In contrast to the clear-sky fluxes, the CREs in the first predicted low sensitivity ensemble demonstrates systematic biases (yellow lines in the middle column of figure 5.7). Whilst mid-latitude and sub-tropical SW biases are generally not systematically different to the standard HadCM3, tropical biases of >25Wm$^{-2}$ exist for the SW-CRE. A typical feature even in CMIP5 GCMs (the generation after HadCM3) is an overly negative net SW-CRE in the tropical regions (∼-15-20Wm$^{-2}$) [Flato, 2013], but such strong biases as seen in the first predicted low sensitivity ensemble would lie outside the range of CMIP5 GCM biases in this region. The LW-CRE shows a compensating heating bias over the same region, resulting in a net CRE within the tropical bias range of the CMIP5 ensemble (∼-25-10Wm$^{-2}$). Compensating SW-CRE and LW-CRE biases are generally expected due to well-understood physical mechanisms in which changes in optically thick cloud are associated with changes in optically thin cloud that have strong LW radiative effects over the tropics [Hartmann *et al.,* 2001].

In an attempt to remove these tropical CRE biases, I construct additional emulators of tropical SW-CRE and LW-CRE. These emulators were constructed and trained in the same way as described in section 5.2 using the averages of the respective component of CRE over 30°N-30°S. The green lines in figure 5.7 show a second iteration of the low sensitivity ensemble with additional constraints imposed upon the emulated control climate SW-CRE and LW-CRE to be within ±5Wm$^{-2}$ of the CERES observations averaged over the 30°N-30°S region. Whilst this additional constraint was generally successful in reducing the tropical CRE biases to within the desired tolerance, an additional consequence of this constraint was the introduction of new biases >10Wm$^{-2}$ in the SW-CRE over the sub-tropical regions. SW-CRE bias is again compensated by an opposing LW-CRE bias, although the simulation of LW-CRE in many of these GCMs is as good as many CMIP5 GCMs over large parts of the globe. The red lines in figure 5.7 show a third iteration of the low sensitivity ensemble, in which SW-CRE and LW-CRE are additionally constrained to have emulated absolute biases from CERES observations of <5Wm$^{-2}$ over the regions

60°S-40°S and 40°N-60°N. Under all three constraints on the emulated radiative fields, the simulated radiative characteristics of this iteration of the low sensitivity ensemble fail to satisfy all the constraints and approximately halves the tropical and subtropical biases present in the first and second iterations of the low sensitivity ensemble respectively. SW-CS fluxes over the Arctic are generally improved in this iteration of the low sensitivity ensemble driven by a better representation of the seasonal cycle of sea-ice (see chapter 6).

### 5.5.2 The climate response of the low sensitivity ensemble

This section analyses the climate response in the low sensitivity ensemble and highlights commonality and diversity within it. Detailed analysis of the underlying physical mechanisms is discussed in the next chapter.

Figure 5.8 shows the temporal evolution of GMST and global-mean precipitation under STEP and RAMP scenarios, for members of the training (grey) and predicted low sensitivity ensemble (blue) across all iterations of the low sensitivity ensemble presented in figure 5.7. The yellow lines show the subset of the predicted low sensitivity ensemble that are found to have an ECS<2K, when calculated through the regression method. The predicted low sensitivity ensemble has clearly been successful in exploring a lower climate response under doubling of $CO_2$ than sampled in the original training ensemble. The thick grey line shows the 10th percentile of the training ensemble and highlights the divergence between the low sensitivity ensemble and the training ensemble with the standard HadCM3 serving as an upper limit to the predicted low sensitivity ensemble response. Many GCMs have global-mean precipitation responses that are still drier than their control climates after 70 years of doubled $CO_2$ integration. Fitting a simple model, based on the global-mean tropospheric energy budget, for equilibrium precipitation changes as a linear function of ECS [Allen & Ingram, 2002], several of the lowest ECS GCMs have

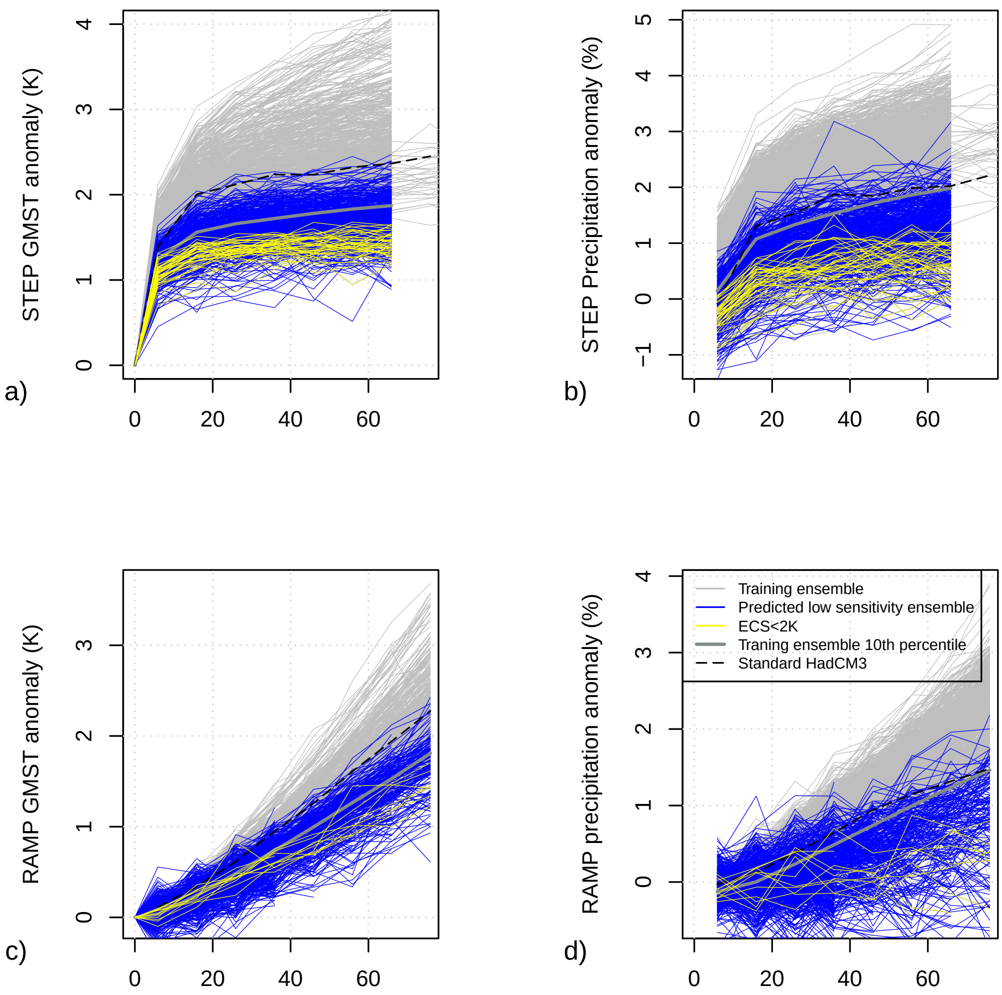


**Figure 5.8** – GMST and global-mean precipitation evolution in the STEP and RAMP experiments for the training ensemble from chapter 3 (grey) and the predicted low sensitivity ensemble (blue). Yellow lines show ensemble members that are found to have ECS<2K. Standard HadCM3 is shown with a dashed black line and the 10th percentile of the training ensemble distribution is shown with a thick grey line.

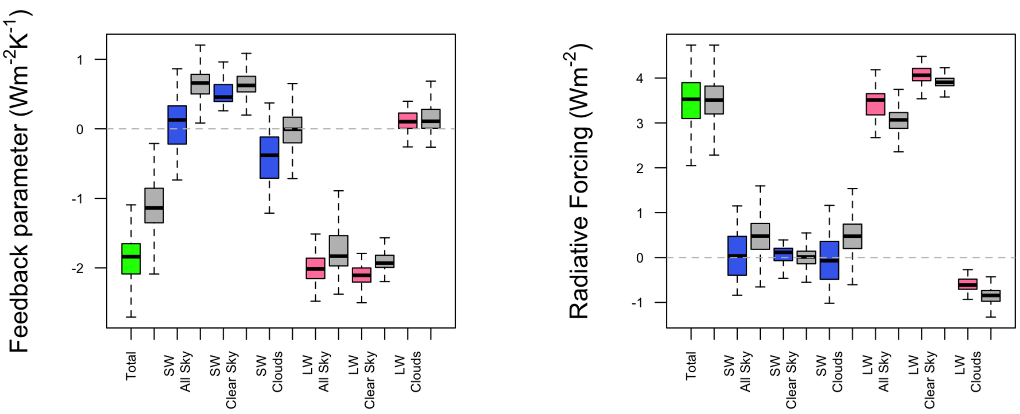


**Figure 5.9** – Components of the global-mean feedback parameter (left) and radiative forcing (right) in ensemble members with ECS<2K (coloured bars). The grey bars represent the corresponding distributions from the training ensemble.

expected global-mean equilibrium precipitations less than 1% wetter than their control climates, but none have low enough sensitivities that expected equilibrium precipitation would be drier than their control simulation in the global-mean.

Across the entire predicted low sensitivity ensemble, 33 GCMs (yellow) are found to have ECS<2K when calculated with the regression-based method of Gregory *et al.* [2004]. For clarity, I will hereafter call these GCMs the 'actual low sensitivity ensemble'. ECS ranges between 1.44-1.99K, with a mean of 1.86K, within these GCMs. Global-mean feedback parameters (and their components) are evaluated across these 33 GCMs in the left hand panel of figure 5.9. The median global-mean feedback parameter for the actual low sensitivity ensemble is -1.8Wm$^{-2}$K$^{-1}$, outside the interquartile range of the training ensemble. This difference is mainly driven by the simulation of more negative SW cloud feedbacks, which are negative in all but 4 of the actual low sensitivity ensemble, in contrast to the training ensemble in which there are roughly equal numbers of GCMs with positive and negative cloud feedbacks. Simulation of LW cloud feedbacks is almost identical (except for a reduction in the maximum positive LW cloud feedbacks) between the two ensembles indicating that Iris-like effects, which would impact the ranges of LW cloud feedbacks simulated in the actual low sensitivity ensemble, isn't likely to be the driving mechanism for low ECS in these GCMs.

Overall radiative forcing estimates are very consistent with the training ensemble and other estimates of the global-mean effective radiative forcing from doubled atmospheric $CO_2$ (e.g. Forster *et al.* [2013]). Clear-sky forcings remain very well constrained in these GCMs and the actual low sensitivity ensemble shows some a more negative fast adjustment in the SW-CRE compared to the training ensemble and conversely slightly less negative fast adjustments in the LW-CRE. The opposite sign of these differences from the training ensemble might be expected due to the physical coupling of LW and SW-CREs in the tropics as previously discussed. As observed in the training ensemble, global-mean forcing and feedback remain correlated (r=-0.73) with each other in these ensemble members.

Figure 5.10 shows the decomposition of cloud feedbacks and forcing by geographical regions. These regions are chosen to match the regional decomposition in Webb *et al.* [2006], based on separating major features of CRE control climates and to avoid sea-ice in mid-latitude boxes. SW-CREs show the greatest variance across the low sensitivity ensemble and are the main drivers of variations in the net feedback. SW-CRE feedbacks show the greatest spread over the low-latitude oceans (LLO). Although many of these GCMs show positive LLO SW-CRE feedbacks, this is compensated by an opposing SW-CRE fast-adjustment which increases SW-CRE cooling of this region. SW-CRE feedbacks and effective radiative forcings are strongly correlated (r=-0.92) over the LLO region resulting in a increase in cooling via the SW-CRE over LLO at the end of the integration period in nearly all GCMs (not shown), indicating that the partitioning between forcing and feedback over these regions may be highly sensitive to internal variability, questioning the physical conclusions that can be drawn from the this separation of the response into feedback and forcing components on regional scales. Non-negligible cooling effects in SW clouds are also generated from the mid-latitude oceans, and higher latitudes. In contrast to SW-CRE, LW-CRE displays significantly less variation across the GCMs with the mean feedbacks being close to zero in all regions. Initial condition ensembles with these parameter sets would help reduce the impact of

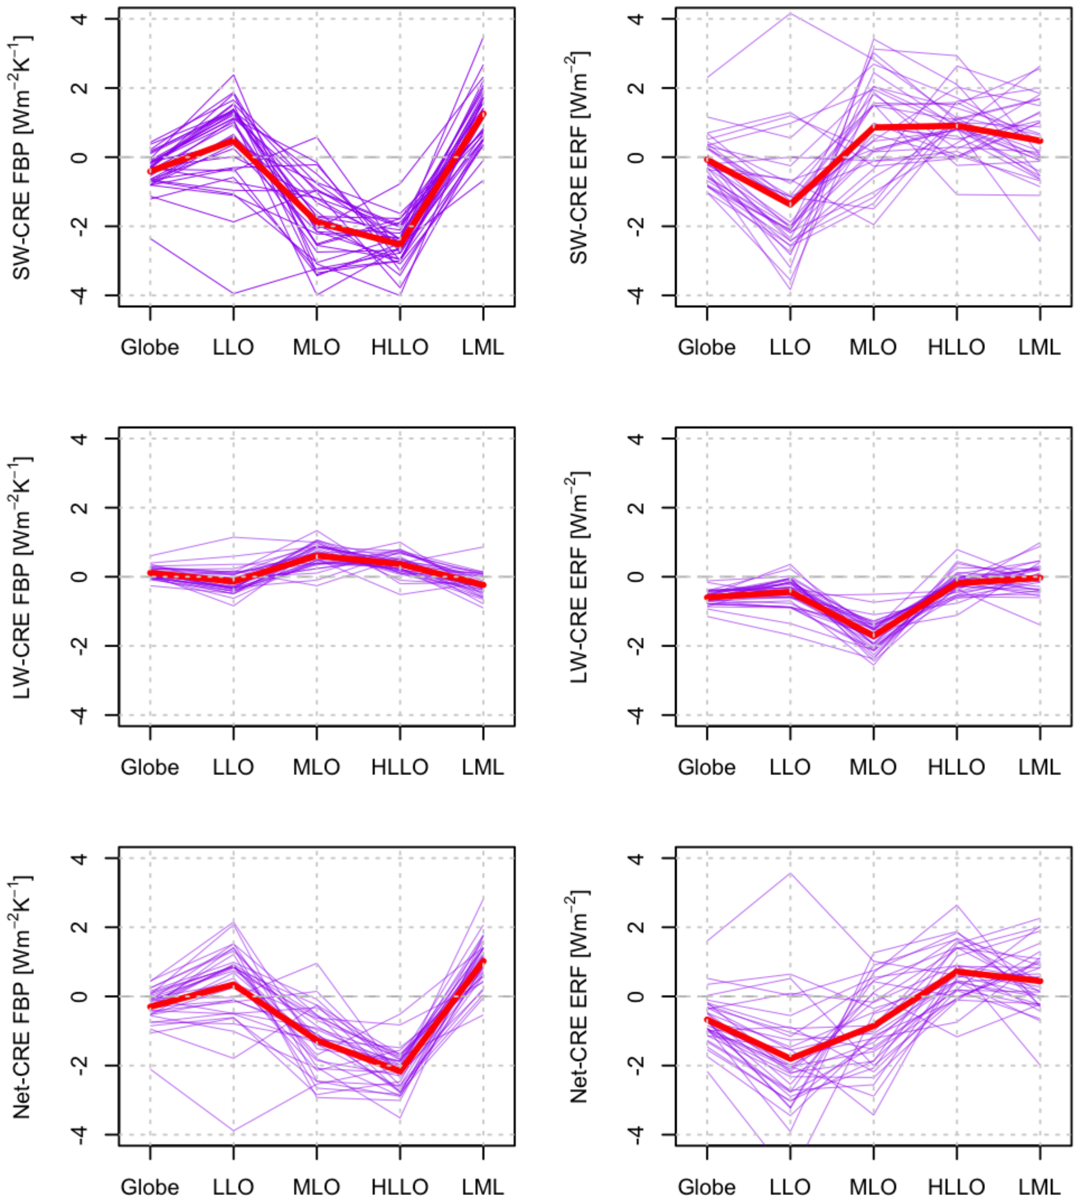


**Figure 5.10** – Decomposition of cloud feedbacks and effective radiative forcing by region for ensemble members with ECS<2K. The regions are defined as follows: LLO - low-latitude ocean (ocean, 30°N-30°S), MLO - mid-latitude ocean (55°N-30°N and 30S-55S), HLLO - high latitude land/ocean (land and ocean; 90°N-55°N and 55°S-90°S) and LML - low/mid-latitude land (land; 55°N-55°S). The thick red line shows the actual low sensitivity ensemble mean with ensemble members shown in purple.

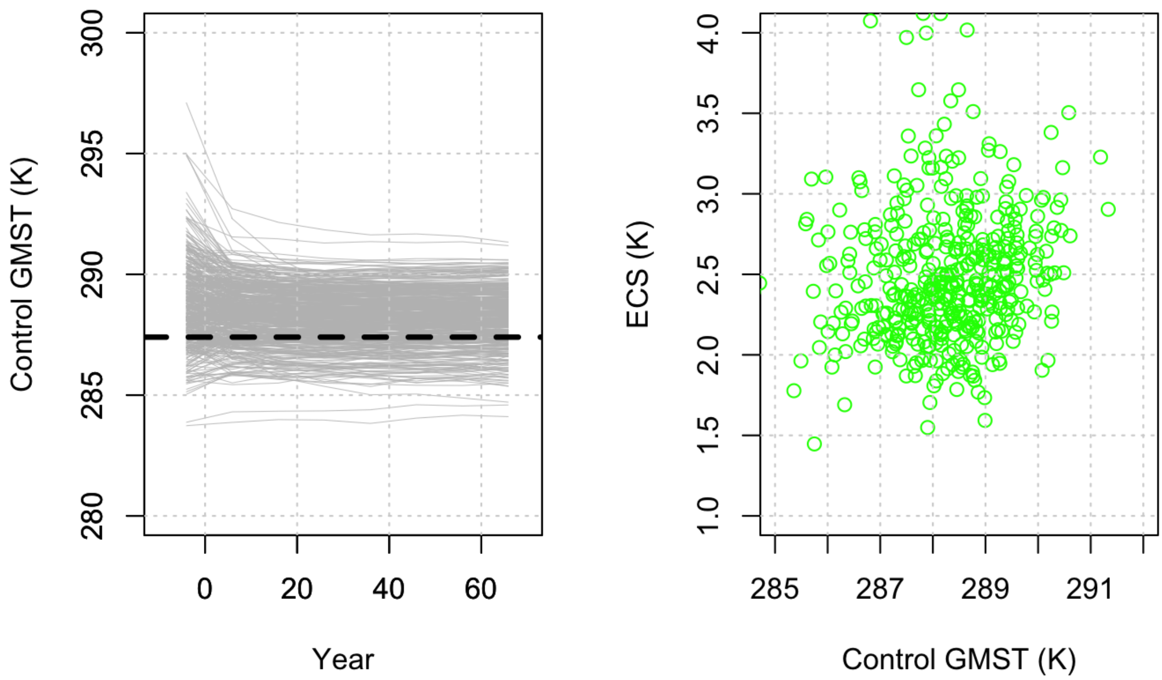


**Figure 5.11** – Left: evolution of control GMST in the predicted low sensitivity ensemble. The thick and dashed black line corresponds to the average GMST from ERA-Interim (1979-1998) [Dee *et al.*, 2011]. Right: scatter plot of control GMST and simulated ECS.

natural variability and increase the confidence in the regional partition between forcing and feedback.

Is ECS simply associated with any large-scale properties of the climate system in the predicted low sensitivity ensemble? Figure 5.11 (left) shows the temporal evolution of the GMST in the control simulations for the predicted low sensitivity ensemble. The first decade of the simulations shown in figure 5.11 displays adjustments in GMST associated with the disequilibrium of the atmosphere and ocean due to the perturbation of the atmospheric parameters. STEP experiments are not commenced until year 10 of the control integration to remove influence of these initial transients on the anomalies. In general control GMST is not a skilful predictor of simulated ECS within the predicted low sensitivity ensemble, as shown by the lack of correlation in the right hand panel of figure 5.11. Figure 5.12

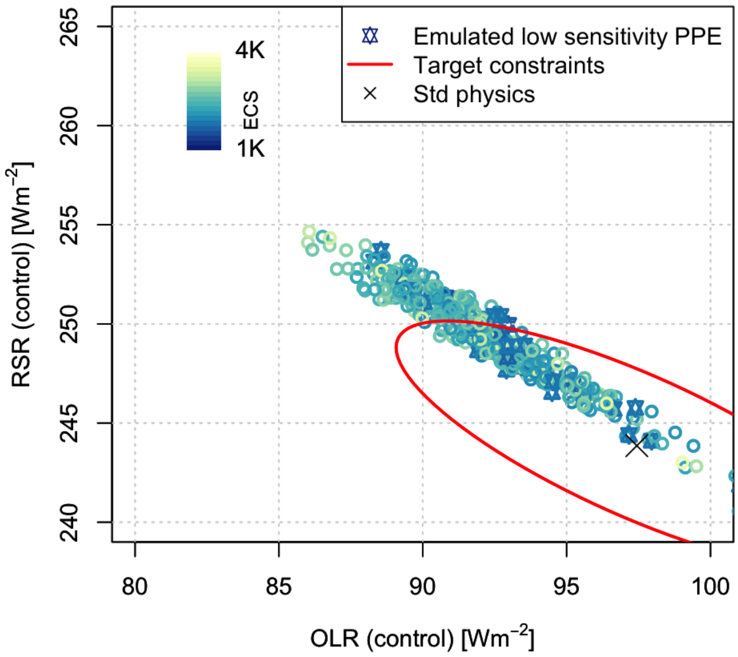


**Figure 5.12** – OLR and RSR evaluated over the final 20 years of control integration, coloured by ECS. The red ellipse shows the 99% confidence region around the standard HadCM3 (black cross) with uncertainty covariance described by equation 5.4.2. Star symbols are used to denote ensemble members with ECS<2K.

shows control OLR and RSR in the predicted low sensitivity ensemble coloured by their simulated ECS values. On average, the highest ECS GCMs lie outside of the targeted constraints (the red ellipse shows the 99% confidence region described by the covariance matrix of equation 5.4.2) and most GCMs with ECS<2K are found closer to the centre of the ellipse (located at the standard HadCM3 OLR and RSR). However, many exceptions to this characterisation can be seen in figure 5.12, indicating that control climate radiative fluxes are not, by themselves, sufficient predictors of the ECS in this ensemble. Whilst some ECS<2K GCMs fall very close to the standard version of the GCM, the global-mean quantities on the axes of figure 5.12 masks the existence of regional CRE biases shown in figure 5.7.

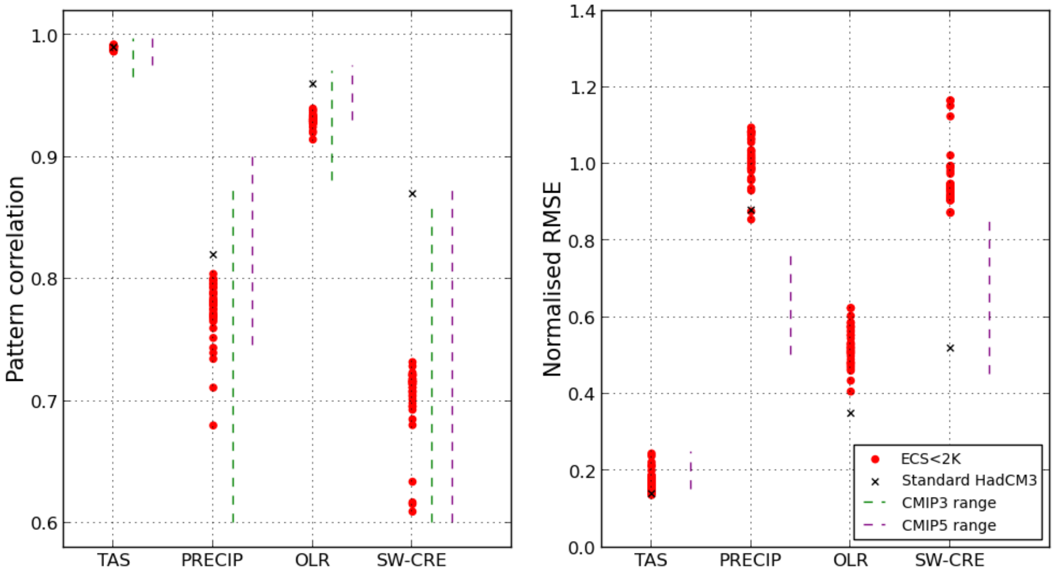


**Figure 5.13** – Left: Pattern correlations relative to observational or reanalysis data [Taylor, 2001] and right: normalised root mean squared error from observations (RMSE - normalised by observational standard deviation). Fields shown are surface air temperature (TAS), precipitation (PRECIP), outgoing LW radiation (OLR) and SW-CRE. Red dots show ECS<2K PPE members. The dashed lines show the CMIP3 (green) and CMIP5 (purple) ranges, where data was obtainable. The standard version of HadCM3 is marked with black crosses.

### 5.5.3 Control climate simulation

There are many metrics of GCM fidelity to observations that can be used to assess GCM skill, but here I choose to focus on two measures that can be used to assess the overall bias magnitude and spatial similarity of GCM climates relative to observations [Taylor, 2001] (figure 5.13). Temperature biases are evaluated relative to the ERA-Interim reanalysis period 1979-1998 [Dee *et al.,* 2011], precipitation relative to the Global Precipitation Climatology Project (GPCP) [Adler *et al.,* 2003] and TOA CRE fields are evaluated relative to the CERES observations [Loeb *et al.,* 2009]. Whilst the pattern correlations for all members of the actual low sensitivity ensemble are reduced from the standard HadCM3 (black crosses in figure 5.13), particularly in SW-CRE, for all spatial fields the range of the low sensitivity ensemble is within the CMIP5 and CMIP3 MMEs. The pattern correlations for the actual low sensitivity ensembles are generally very close to the mean of the CMIP3 (mean: 0.72, range: 0.60-0.86) and CMIP5 (mean: 0.73, range: 0.60-0.88) ensembles, indicating comparable skill in the spatial representations of climatological features (often indicative of simulating some elements of the large-scale circulation correctly) to most CMIP3/5 GCMs. In terms of absolute errors (RMSEs; right hand panel of figure 5.13), the actual low sensitivity ensemble is only marginally degraded from the standard HadCM3 for 3 out of the 4 fields shown (with some ensemble members showing improved precipitation and surface temperature climatologies). However normalised RMSEs for SW-CRE are substantially degraded from the standard HadCM3. As HadCM3 was one of the better GCMs on this metric, even in the CMIP5 MME, despite this substantial degradation, the majority of the actual low sensitivity ensemble members lie just outside of the range of the next-generation CMIP5 GCMs.

These globally-averaged, but none-the-less important, metrics of model skill indicate that many low response GCMs would perform remarkably well when compared to state-of-the-art GCMs from the CMIP5 ensemble. CMIP5 GCMs have both higher spatial resolutions and have undergone thorough tuning processes to improve the fidelity of their climate simulations. As the standard HadCM3 performs very well relative to other CMIP3 and CMIP5 models across nearly all of these variables, even degraded performance from the standard HadCM3, would not make many of these low ECS GCMs complete outliers if they were included in the multi-model distribution. As SW-CRE effective radiative forcing and feedbacks are important for the simulation of low ECS in the low climate response ensemble, it may be that features of the SW-CRE field (that is degraded at the global scale) are essential to the simulation of these forcings and feedbacks. A more detailed and in-depth examination of the simulation of cloud properties will be examined in the next chapter.

## 5.6 Discussion and conclusion

This chapter set out to explore the lower bound of climate response, motivated in part by the possibility that existing GCM ensembles may be insufficiently diverse representations of climate response space (see chapter 4). In contrast to other recent studies, which have sought to create low response GCMs through the explicit introduction of a specific mechanism [Mauritsen & Stevens, 2015], I have instead attempted to find low response GCMs within the parameter space of the HadCM3 GCM without introducing any new modifications to the model.

To do this, a statistical tool, known as an emulator, was constructed to emulate the climate response space of the HadCM3 PPE. This emulator, described in detail in the first half of this chapter, generated a prediction for a set of points in parameter space that were expected to have realistic TOA radiative fluxes in their control climates as well as ECS<2K. A common feature of these GCMs, as well as GCMs in general, was the existence of biases in the tropical and subtropical CREs.

From the predicted low sensitivity ensemble, 33 members that completed the integration were found to have ECS of less than 2K. These GCMs have an average ECS of 1.86K and a low value of 1.44K. The low ECS in these GCMs is largely driven by negative SW global cloud feedback, which is particularly apparent over the tropical oceans, indicating that the low climate response in these GCMs is not predominantly driven by an Iris effect which would be manifest in the LW cloud feedbacks of these GCMs. At a global-mean level, the low ECS results primarily from more negative SW cloud feedbacks, although fast adjustments are also shown to have an important effect in certain regions when the climate response is evaluated using a regression-based approach (although the physicality of such fast-adjustments is unclear due to strong correlations between forcing and feedbacks). A detailed examination of the physical mechanisms governing the climate response in the low sensitivity ensemble is left until the next chapter.

The assessment of the skill of these GCMs can be used to consider the relevance of these GCMs for understanding climate sensitivity in the real climate system. In relation to this, I here highlight several points of consideration. Whilst the perturbed GCMs discussed in this chapter generally show reduced skill relative to the standard HadCM3, these GCMs have received less tuning than the 'flagship' GCMs in the CMIP5 ensemble, which receive many man-months of fine-tuning to improve their climate simulation before being finalised [Mauritsen *et al.,* 2012]. Additional investment in tuning these low ECS perturbed GCMs could reduce many of the control climate biases that are present in these GCMs without substantially altering their climate sensitivity.

When comparing these perturbed GCMs to the CMIP5 MME, the un-tuned low ECS GCMs would lie within the CMIP5 ensemble on nearly all metrics and would lie just outside of the current CMIP5 range for SW-CRE RMSEs. As highlighted in the text previously, CMIP5 GCMs are the next generation of GCMs after HadCM3 and generally exhibit increased horizontal and vertical resolution as well as updated and improved parametrisation schemes. I therefore argue that as these low resolution and crudely-tuned perturbed GCMs lie only marginally outside the CMIP5 range there is insufficient evidence to convincingly rule out the possibility of a low climate sensitivity (between 1.5K and 2K) based on this experiment. The results in this chapter suggest that it may be possible to produce plausible looking state-of-the-art GCMs with low climate sensitivities if GCM modelling centres set out with this as an explicit goal and were similarly able to fully explore the parameter space of plausible climate response in their GCM framework.

Potentially important additional degrees of freedom, that haven't been explored within this ensemble, exist with HadCM3. One of these degrees of freedom is the use of the HadCM3 anvil perturbation scheme. This scheme allows for an enhancement of the cloud area as seen by the GCM radiation scheme, to mimic the radiative impact of convective anvils in deep convective regions. This perturbation could allow additional degrees of freedom related to tropical cloud properties and may interact with perturbations to other GCM parameters leading to potentially important effects on both the ECS and the control climate CRE. Including perturbations in the GCM anvil factor could therefore offer new ways of rectifying remaining CRE biases whilst preserving low ECS. Another potentially important parameter that is currently not perturbed with HadCM3 is the parameter controlling the partitioning of condensed water between liquid and solid phase at a specified temperature. This parameter is involved in determining the properties of mixed-phase clouds and the average cloud albedo (cloud ice is less reflective than liquid cloud water). It has been shown that changes to this parameter can affect the phase-change feedback under warming (a rising melting level reduces the ice content of clouds all else being equal), particularly over high and mid-latitudes oceans in other GCMs [McCoy *et al.,* 2016]. Perturbations to this parameter may offer additional mechanisms for achieving low climate response different from the dominant low-latitude ocean signal in these PPEs.

Only 6 of the ECS<2K GCMs have currently completed and returned data for the RAMP experiment from the climate*prediction*.net computing infrastructure (although more of the 33 members of the actual low sensitivity ensemble will be

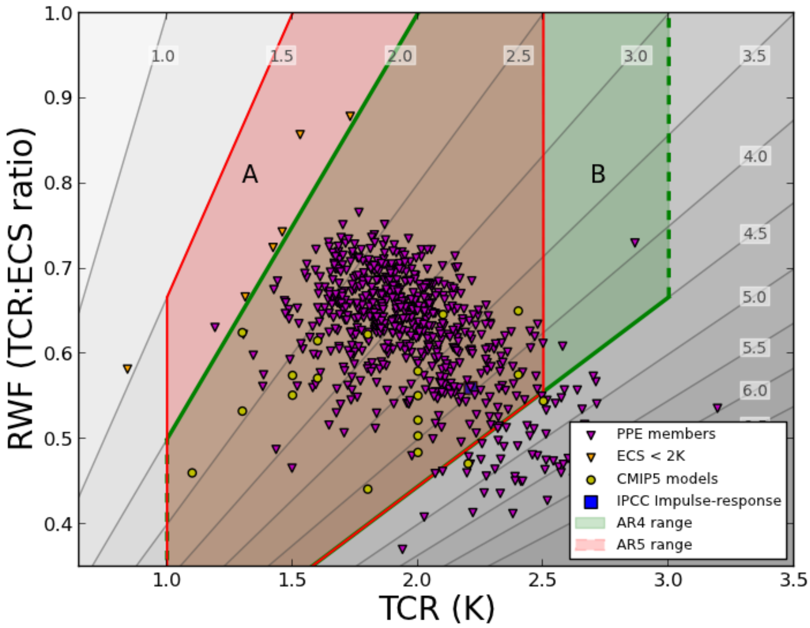


**Figure 5.14** – As in figure 4.8, but with ECS<2K GCMs that also return a TCR value added as orange triangles.

expected to return RAMP experiment data in the future). These points are added to the TCR/RWF plot as shown in figure 4.8. TCR values are typically much better constrained than ECS and despite the changes in the joint TCR/RWF distribution through the additional low response GCMs, new low TCR ranges have not in general accompanied the decrease in simulated ECS (RWF has increased instead). The lowest ECS member of the actual low sensitivity ensemble (ECS=1.44K) had a TCR=0.84K, but only two other members of either the actual or predicted low sensitivity ensemble have TCR<1.0K (both located at 0.95K). The maximum TCR in the predicted low sensitivity ensemble is 2.3K and 97% of the distribution lies above a TCR of 1K. These results again suggest that that TCR is a much better constrained measure of the climate response than ECS, including from GCM ensembles, as the unprecedented low ECS GCMs in the actual low sensitivity ensemble generally have TCRs within GCM and observationally-derived uncertainty ranges (such as the AR5 *likely* range). As TCRs of as low as 1.0K can be found to be consistent with the historical period when uncertainty in the total radiative forcing is taken into account, it would be expected that, unless non-$CO_2$ forcing has efficacies significantly different from unity, that the actual low sensitivity ensemble would perform well in simulations of the GMST trend over the observed period.

CHAPTER 6

# Mechanisms for low climate response

Chapter 5 described a PPE-based search for GCMs that have ECS<2K. Using a statistical emulator, an ensemble of low climate response GCMs was created, a number of which had ECS<2K. Low ECS across the ensemble was largely driven by changes in SW-CRE, particularly over low-latitude oceans, but physical mechanisms for these responses where not analysed in chapter 5.

This chapter investigates the physical mechanisms underlying variations in climate response across the low sensitivity ensemble. Section 6.1 analyses a representative subset of the predicted low sensitivity ensemble, aiming to understand features of the control climate simulation. Features of these GCMs' control climates are contrasted with the standard HadCM3 as well as various observational and reanalysis datasets (section 6.2) to assess GCM accuracy. Section 6.3 focuses on features of the climate response and decomposes the climate response into contributions from different cloud regimes over low-latitude oceans. Finally, section 6.4 analyses potential physical mechanisms that may underlie the cloud response seen in these

GCMs. This work, along with a shortened summary of the results described in chapter 5 is being prepared for submission to *Climate Dynamics*.

## 6.1 The local ensemble

Storage limitations of the climate*prediction*.net distributed computing project require spatial fields of geophysical variables from very large ensembles to be output as decadal averages. 3D spatial data with a greater time resolution can be of more value than 2D lower time-resolution data when investigating physical mechanisms that may drive the low climate response to $CO_2$.

To investigate physical mechanisms of climate response across the predicted low sensitivity ensemble, I re-simulate the STEP experiment for a small selection (8 GCMs) of the predicted low sensitivity PPE in a local (non-distributed) computing environment therefore allowing 3D fields at monthly-mean time resolution to be stored. Of these 8 GCMs, 1 is chosen to be the standard HadCM3, and the other 7 are selected to cover the range of GMST response across the predicted low sensitivity ensemble including GCMs from the multiple iterations of regional CRE biases described in section 5.5.1. I choose not to focus solely on the GCMs with ECS<2K, but instead on a range of climate responses from across the ensemble in order to better understand the physical drivers. This set of 8 GCMs will be referred to as the *local ensemble* hereafter, in order to distinguish this set of simulations from the *low sensitivity* ensembles of chapter 5. These GCMs form a representative set of the diversity in climate response and control climate simulation across the predicted low sensitivity ensemble.

Figure 6.1 shows the evolution of the GMST anomaly in all 8 GCMs as a function of time after $CO_2$ doubling. I associate each GCM with a characteristic letter and colour that is consistently used to refer to each GCM in the figures throughout this chapter. Section 4.2 showed the TCR to be the most relevant metric of the GMST

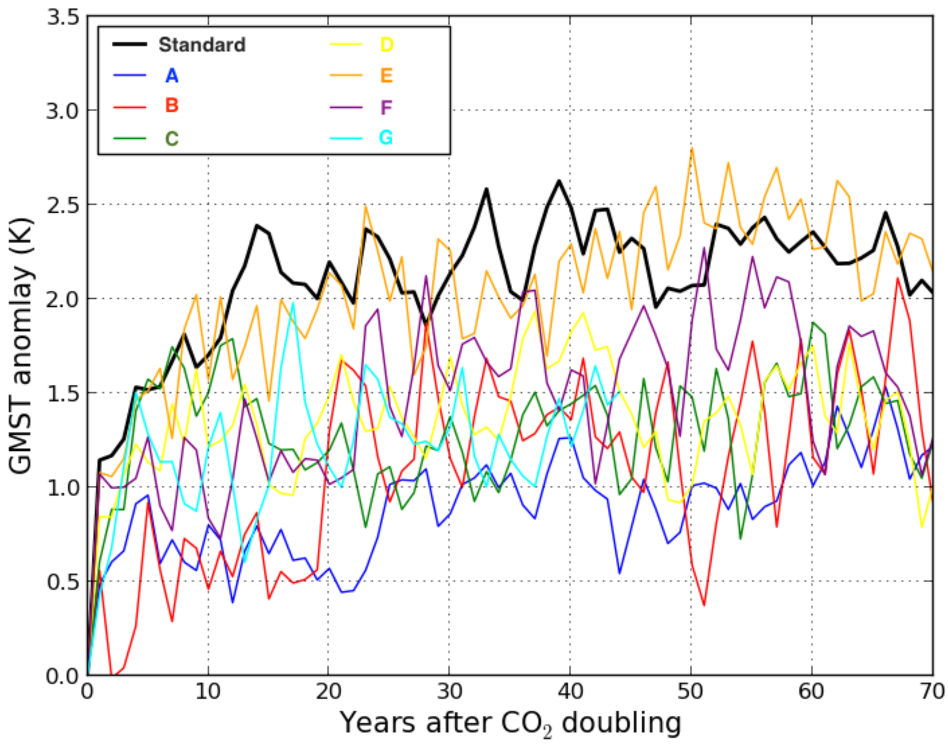


**Figure 6.1** – The evolution of GMST in the 8 GCMs considered in this chapter, given as annual averages for years after instantaneous doubling of atmospheric $CO_2$ concentrations. All GCMs were integrated for 70 years following doubling of $CO_2$ except for one member ('G' - light blue) whose integration failed before completion of the 70 years. The legend shows a letter that is used to denote each member of the local ensemble. In all figures appearing in this chapter, the same letter and colouring is consistently used to identify the GCMs.

| GCM | TCR (K) | ECS(K) |
|---|---|---|
| Standard | 2.11 | 3.50 |
| A | 0.91 | 1.74 |
| B | 1.13 | 2.20 |
| C | 1.32 | 1.90 |
| D | 1.36 | 1.88 |
| E | 2.06 | 3.15 |
| F | 1.51 | 2.40 |
| G | NA | 1.78 |

**Table 6.1** – TCR and ECS values for each of the GCMs in the local ensemble analysed in this chapter. ECS is here estimated via total least squares regression with 5-year averages, to help reduce the impact of the low signal-to-noise ratio in the low climate response members.

climate response to radiative forcing for GMST changes under policy-relevant future emissions scenarios, therefore, whilst I focus the analysis on STEP experiments where concentrations are doubled and then held constant (to isolate the response to an invariant radiative forcing), I choose to numerically summarise the overall difference in response between these GCMs via their TCR (derived through the linear step-response method of Good *et al.* [2011], as discussed in section 3.4) and ECS, which are listed in table 6.1. Hardware failures prevented the 'G' local ensemble member from completing its $CO_2$ forced integration and therefore a TCR value is not available. However I continue to show the control climate simulation and normalised climate response patterns throughout this chapter, to inform analysis of physical mechanisms.

## 6.2 Fidelity of pre-industrial control climates

Figure 5.13 assessed GCM fidelity in the simulated pre-industrial control state against observational data for low sensitivity ensemble members with ECS<2K. However, due to constraints on the visualisation of spatial data from large distributed ensembles, the analysis only focused on globally-averaged metrics, which are simple to represent across the ensemble. In particular, 'emergent constraints', in which a property of the climate response is associated with a property of the control climate [Allen & Ingram, 2002; Rowlands *et al.,* 2012; Sherwood *et al.,* 2014b; Tian, 2015] may be found on non-global scales [Tian, 2015]. In this section I analyse spatial visualisations of GCM fidelity for the smaller local ensemble.

GCMs displaying smaller simulation errors for particular fields are generally assumed to better capture the physical processes underlying that simulation, and may be given a higher weight in assessments of the response of this process to enhanced $CO_2$ concentrations [Murphy *et al.,* 2009]. Figure 5.13 showed that overall GCM fidelity was worse than the standard HadCM3 in PPE members with ECS<2K. It is expected that perturbed GCMs would in general have reduced skill when compared to the standard 'flagship' HadCM3, which received many man-years of tuning to remove biases. However, it may be that low ECS could be associated with the reduction (or enhancement) of particular biases in these GCMs and therefore examining control climate biases may offer an indication of which process are important for the simulated climate response.

Figure 6.2 shows the near surface air temperature for the local ensemble, expressed as errors relative to the earliest available 20 year ERA-Interim reanalysis period (1979-1998) [Dee *et al.,* 2011]. Temperature observations are not globally complete, particularly in the early historical record and the Southern Hemisphere [Morice *et al.,* 2012]. Hence, I use reanalysis data to compare full spatial fields with the global GCM output. I contrast the earliest available reanalysis period with the GCM control state, defined as a temporal average of the final 20 years of GCM integration. Pre-industrial conditions in the GCM are defined as 1850 values of atmospheric concentrations and forcings. Whilst the external forcing on the climate system is not identical between the GCM control and the reanalysis period, local GCM biases are typically much larger in magnitude than the local climate change signal arising between 1850 and the end of the twentieth century. Therefore, a direct comparison between the GCMs and the reanalysis is a useful indicator of GCM skill in representing the surface climate.

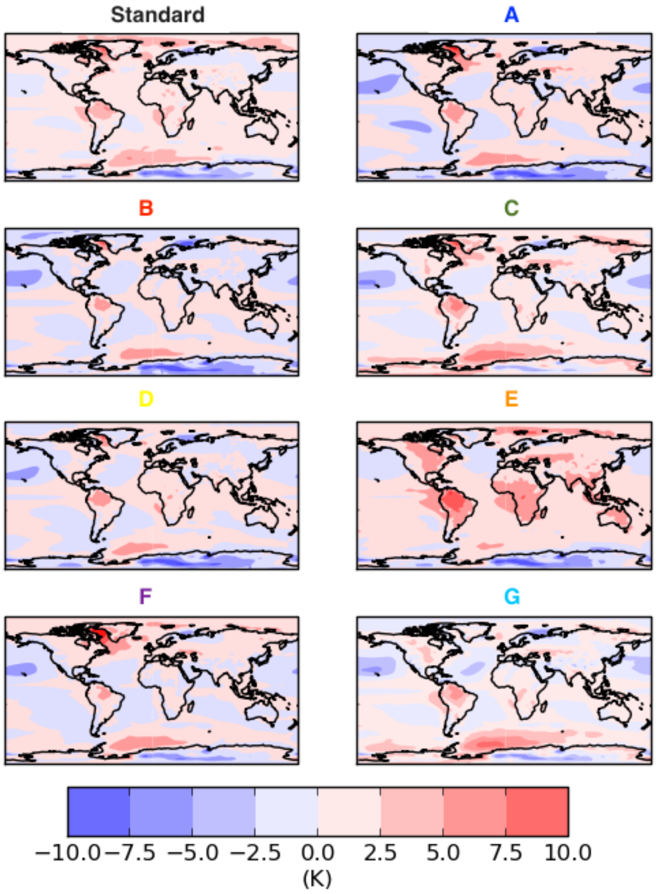


**Figure 6.2** – Biases of GCM control climate near (1.5m) surface air temperature from the ERA-Interim reanalysis dataset of (2m) surface air temperature.

Across the local ensemble, typical patterns of surface air temperature biases are roughly consistent with the biases in the standard HadCM3. In particular, substantial biases are observed over Antarctica and the Southern Ocean. These biases are co-located with the under-simulation of sea-ice cover in the GCMs, particularly east of the Antarctic Peninsula and in the austral winter, compared to the HadISST reanalysis (not shown). Similarly, the warm biases in the Arctic are co-located with under-simulated sea-ice. The standard HadCM3 displays a particularly strong temperature bias in the Arctic, accompanied by a very low simulated sea-ice minimum in the late boreal summer (see figure 6.3). This summer sea-ice bias is not present in the original Met Office HadCM3 variant [Gordon *et al.,* 2000], and was discovered as part of this work (see chapter 3). This bias can be traceable to the inclusion of an interactive sulphur-cycle in the GCM, ultimately leading to

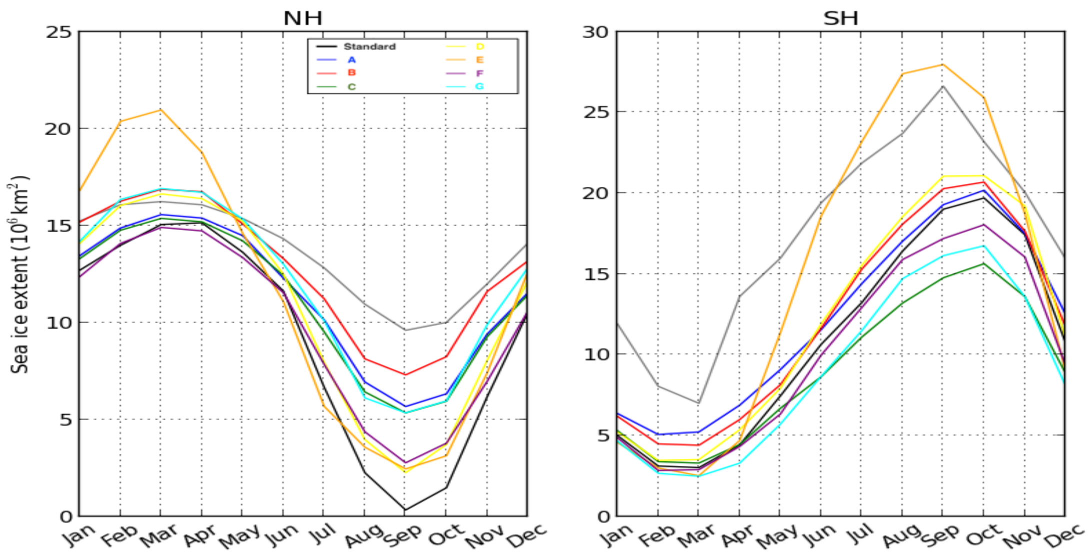


**Figure 6.3** – Annual cycle of the sea-ice extent (area of grid-boxes with sea-ice coverage more than 15%). HadISST [Rayner *et al.*, 2003] 1870-1889 period (grey) and the local ensemble (colours), for the Northern (left) and Southern (right) hemisphere.

warm biases existing through the upper 500m of the GCM ocean. This has been a thus-far undetected GCM bias in all standard physics versions of HadCM3 used for perturbed physics studies. Further work is currently being undertaken to remove this bias, however a comparison with a UK Met Office HadCM3 control integration (not shown) shows that this bias affects the GCMs only in the high latitudes with tropical and extra-tropical regions unaffected. Additionally, I see minimal evidence of any effect of this bias on the climate response of the standard HadCM3, and therefore believe that this bias has a minimal impact on the comparisons of climate response that are undertaken later in this chapter.

Some low response GCMs reduce biases present in the standard HadCM3 over tropical land, in particular over northern continental South America and Central Africa. These biases are still common in recent GCMs and are partly driven by a lack of sufficient upwelling off western coasts [Flato, 2013]. The GCM 'E', which has a GMST response closest to the standard HadCM3 (figure 6.1), displays even more pronounced warm biases over these upwelling regions and over most land masses.

Precipitation, which depends strongly on the interaction of resolved GCM dynamics with several different sub-gridscale parametrisation schemes, is expected to be more sensitive to parameter perturbations than temperature fields. Figure 6.4 shows the biases of the GCMs' precipitation climatology relative to the Global Precipitation Climatology Project (GPCP) [Adler *et al.,* 2003]. The climatology in the GPCP dataset is defined as the average across the period 1981-2010 and is a combination of rain gauge data, satellite and sounding observations to form a spatially complete precipitation history over both land and ocean.

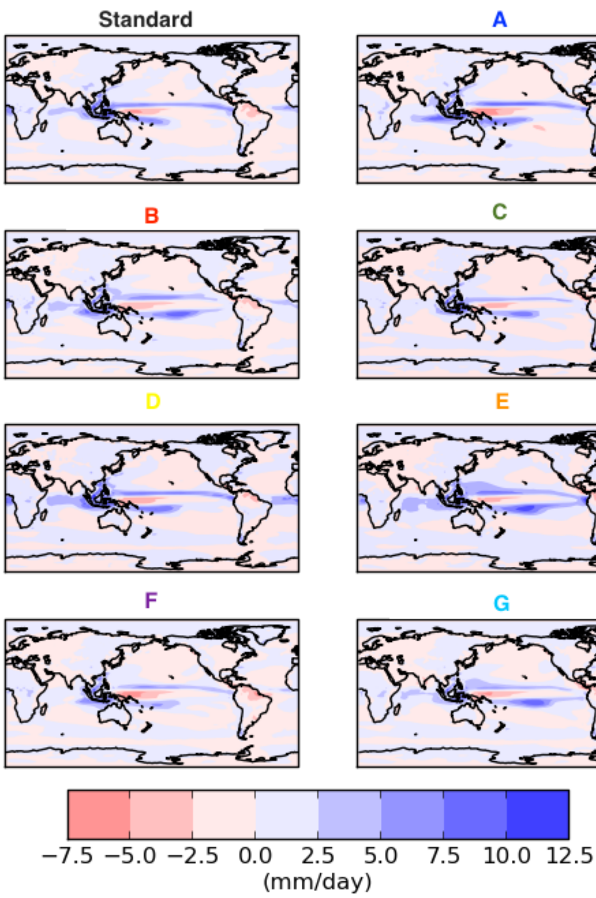


**Figure 6.4** – Biases of GCM climatological total precipitation flux from the GPCP Adler *et al.,* 2003 observational dataset. The final twenty years available of the control integration is used to calculate the GCM climatology which is compared to the earliest available 20 year period in the GPCP dataset (1981-2010).

The spatial patterns of precipitation bias in the perturbed GCMs are similar to the standard HadCM3 over most of the tropics. All these GCMs have too much precipitation in much of the intertropical convergence zone (ITCZ), although ITCZ biases are reduced in some perturbed GCMs over the Atlantic ('F' and 'G'). A feature of many of the GCMs is an exaggeration of the 'double ITCZ' bias already present in the standard HadCM3 and many other GCMs [Tian, 2015]. This spuriously splits the ITCZ over the central Pacific Ocean into two separated branches with an overly dry region between them, unlike the observed continuous ITCZ across the Pacific Ocean. Tian [2015] argued that the magnitude of the 'double ITCZ' bias (there defined as the amount of annually averaged precipitation in a box covering 30°S-0° and 100°W-150°W) is a good predictor of the ECS. However this bias is not significantly or strongly correlated (r≈0.15) with ECS/TCR in the local ensemble. Many of the perturbed GCMs exacerbate the common GCM bias of a too-zonally oriented South-Pacific Convergence Zone that already exists in the standard HadCM3. Equatorial precipitation in the Western Pacific is too low in all GCMs, associated with an overly extensive equatorial cold tongue [Collins *et al.,* 2011].

Over the oceans, the low climate response GCMs are typically cooler than both reanalysis and the standard HadCM3 (figure 6.2). Cool biases can be seen over subsidence regimes across large parts of the tropical and sub-tropical oceans, suggesting differences in control climate oceanic cloud simulation. Several perturbed GCMs show this particularly in a subsidence region of the North Pacific. Figure 6.5 (where CRE biases are given relative to the CERES dataset [Loeb *et al.,* 2009]) shows that this is driven by a strong negative SW-CRE (cooling) bias. The anomalously strong SW-CRE cooling in this region is accompanied by a modest increase in cloud fraction (where the standard HadCM3 had too little cloud - as shown in figure 6.14), suggesting that the SW-CRE bias is likely obtained through both higher cloud fraction and cloud albedo in this region.

Figure 6.5, shows other pronounced features of the SW-CRE bias. In all perturbed GCMs, except 'D' and 'E', patches of over negative SW-CRE across large regions of tropical and sub-tropical oceans dominate the pattern of biases of net CRE. These are associated with the simulation of shallow cumulus/stratocumulus clouds

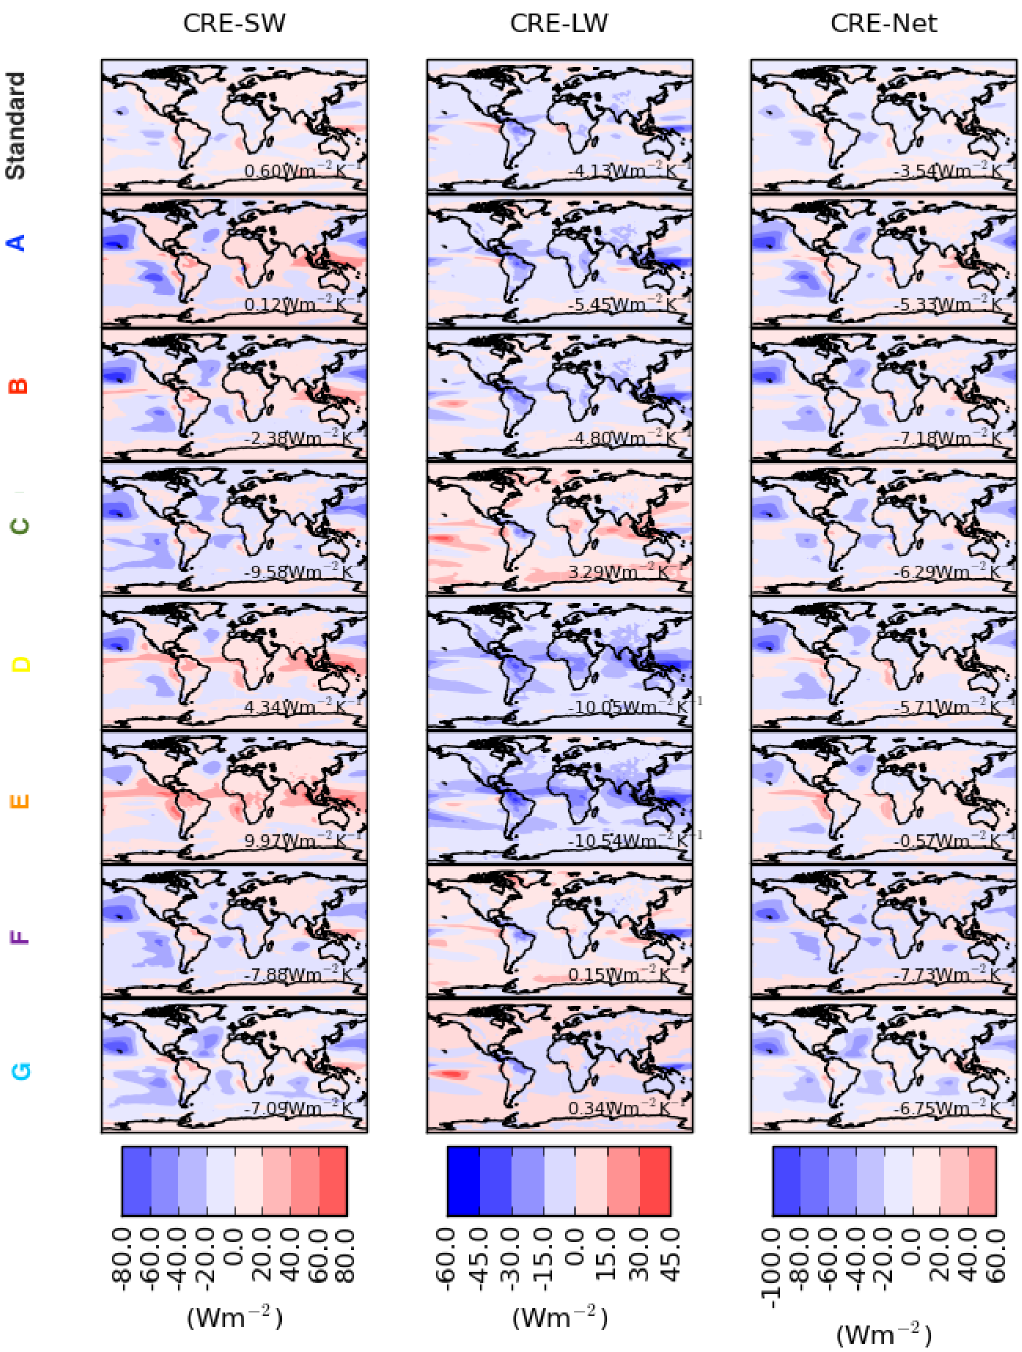


**Figure 6.5** – Annual-average biases in CRE. Biases are calculated using the climatology from the CERES dataset [Loeb *et al.*, 2009]. The numerical value given in each panel shows the global-mean bias.

over subsidence regions and will be explored further later in this section. SW-CRE errors are very strongly correlated with surface air temperature errors ($r = 0.93$) in these regions and across the tropics. In contrast, patterns of LW-CRE errors are very different between GCMs. Two GCMs (again 'D' and 'E') display consistent negative biases of LW-CRE across most of the globe (and particularly over regions of deep convection), but others show both positive and negative LW-CRE biases. The GCMs 'D' and 'E' display similar patterns of CRE bias that are distinct from the other perturbed members. These GCMs are members of the second iteration of the predicted low sensitivity ensemble, all of which have strong negative biases of LW-CRE in the tropics which are compensated by positive SW-CRE biases to give approximate energetic balance, and as might be expected through the physically understood coupling of SW and LW-CREs in the tropics [Hartmann *et al.,* 2001].

Engström *et al.* [2015] provide a new way of comparing the simulation of cloud radiative properties against the observations. They propose that observational evidence supports an exponential relationship (black line in figure 6.6) between oceanic monthly-mean grid-box albedo and cloud fraction in the tropics and sub-tropics (60°N-60°S). This implies that cloud optical depth increases with cloud cover. They show that CMIP5 GCMs generally fail to capture this relationship.

Figure 6.6 shows the 2D-histogram of grid-box albedo and cloud fraction for the local ensemble. Whilst there is a hint of non-linearity in the standard HadCM3 at high cloud fraction and high albedo (or more accurately, evidence of a second distinct branch of the distribution), over most of the range of cloud fraction the relationship is broadly linear. The perturbed GCMs display several features absent from the standard HadCM3. Low sensitivity GCMs typically contain a greater density of grid-boxes at higher cloud fraction and higher albedo than the standard HadCM3, more like the observational histogram (which contains cloud fractions continuously up to 1.0, see Engström *et al.* [2015] figure 1), suggesting that a common error in GCMs of too little but too reflective cloud over sub-tropical and tropical oceans [Nam *et al.,* 2012] may be at least partly alleviated (at least on the

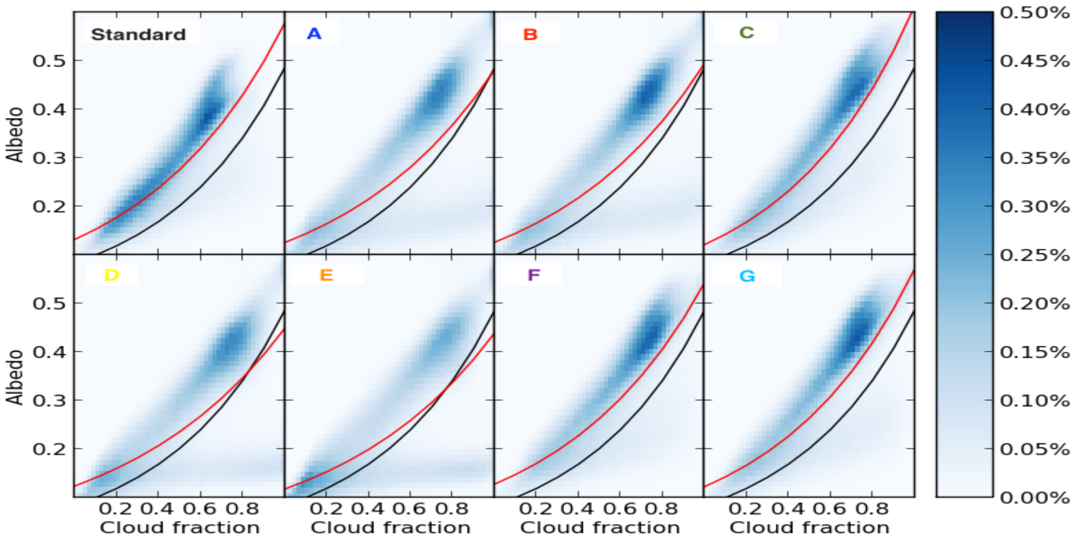


**Figure 6.6** – Histograms of oceanic grid-box cloud fraction and albedo in the monthly-mean between 60°N and 60°S, following Engström *et al.* [2015]. The colouring represents the fraction of total counts across the final twenty years that fall within each bin. The black line represents an exponential fit between grid-box cloud fraction and total albedo in the CERES satellite data sets [Loeb *et al.*, 2009]. The red line shows an ordinary least squares exponential fit to the GCM data.

too little front) in the low sensitivity GCMs.

Additionally, in ‘A’, ‘B’, ‘D’ and ‘E’ evidence can be seen of a separate low albedo branch extending across the cloud fraction range, which is not present in the CERES observations. Approximately constant grid-box albedo over a wide range of cloud fractions (as seen in the low albedo branch) indicates a decreasing average cloud optical depth (decreasing cloud thickness) at high cloud fractions.

At least two local ensemble members seem to exhibit two modes of cloud fraction/albedo combinations in the main branch (‘D’ and ‘E’). A mode exists at high cloud fraction/ high albedo as well as one at low cloud fraction and low albedo, unlike standard HadCM3, where the density is more uniformly spread out along the quasi-linear relationship between cloud fraction and albedo. Exponential fits to the GCM data (red lines) lie above the observational fit (black lines) in all cases. These fits are biased by the low albedo branches and therefore do not pass through the bulk of the histogram density unlike for the observational fit (not shown). Although visual representations of biases are sensitive to the colouring of histogram

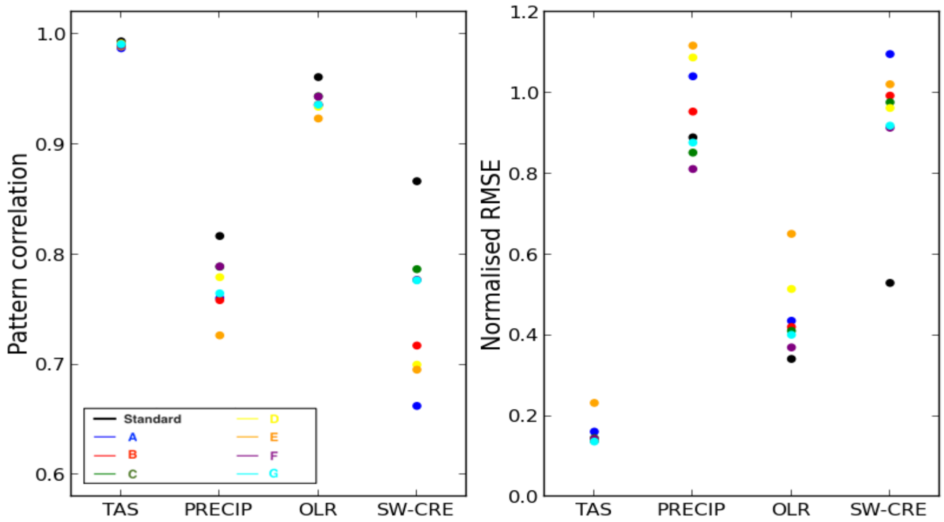


**Figure 6.7** – The left hand panel shows pattern correlations [Taylor, 2001] for surface air temperature (TAS), precipitation (PRECIP), outgoing long-wave radiation (OLR) and short-wave cloud radiative effect (SW-CRE). The right hand panel shows normalised root mean square errors (RMSE) for the same fields as defined in figure 5.13.

density, when assessed through this new framework, fundamental differences in control climate cloud representation are revealed in the perturbed parameter GCMs. However, nearly all CMIP5 GCMs also display histograms densities displaced to the high albedo side of the observational fit (figure 2 of Engström *et al.* [2015]), showing consistent 'too little and too bright' biases across GCMs [Nam *et al.*, 2012]. As many CMIP5 GCMs also display multiple branches in their histogram of these two quantities, the existence of multiple modes and branches within these histograms does not necessarily mean that these perturbed GCMs are performing poorly compared to others in the MME.

As a synthesis of the control climate analysis presented here, figure 6.7 shows the pattern correlations and normalised root mean square error (RMSE), defined as in figure 5.13, for several climatological variables. In all fields (except for RMSE precipitation, which is improved in three of the perturbed GCMs) the standard HadCM3 shows the greatest fidelity to the spatial structure of the observations (pattern correlation) and smallest normalised error (RMSE), with the perturbed variants showing reduced fidelity. In particular, SW-CRE is simulated significantly worse in all the perturbed GCMs, with RMSE 40% lower in the standard HadCM3, despite the local ensemble GCMs having both signs of the global-mean SW-CRE bias in figure 6.5.

As low climate responses in these GCMs are achieved largely through low cloud SW-CRE effects (see next section), the degradation of this field in the control climate relative to the standard HadCM3, raises questions as to the realism of low climate responses. However, it is important to compare the performance of these GCMs to the MME range of CMIP5 and CMIP3. If these GCMs were considered as part of the CMIP5 ensemble, the SW-CRE would lie just outside of the CMIP5 distribution for RMSE (max CMIP5 normalised RMSE in SW-CRE is $\approx$0.9 [Zhao *et al.,* 2016]) and well within the ensemble for pattern correlation. It must be remembered that the HadCM3 is a CMIP3 generation GCM and has lower horizontal and vertical resolution than the CMIP5 GCMs. A comparison solely with CMIP5 GCMs still indicates control climate fidelity has not been decreased in this ensemble to such an extent that these GCMs could clearly be distinguished as complete outliers from the CMIP5 distribution, even for SW-CRE, the worst performing field. The 'flagship' GCMs in CMIP5 have also received dedicated tuning to improve the simulation quality, which hasn't been undertaken (beyond predicted satisfaction of the imposed constrains from the emulator) for the local ensemble and may offer ways to improve the simulations quality further without significantly affecting the ECS.

## 6.3 Climate response in the local ensemble

Having analysed the control climate simulations in the previous section, this section focuses on features of the climate response under a doubling of $CO_2$ in members of the local ensemble and highlights similarities and differences amongst them.

Across the entire low sensitivity ensemble, described in chapter 5, SW-CRE feedbacks, particularly over LLO, largely determined the spread in global climate response. Figure 5.10 showed this was true for GCMs with ECS<2K when evaluated using the linear regression method of Gregory *et al.* [2004]. Understanding the physicality of the partitioning between tropospheric fast adjustment and feedbacks, evaluated through regression-based approaches, can be difficult on a non-global scale, particularly for GCMs with a low climate response and therefore a relatively small signal-to-noise ratio. Local fast adjustments and feedback components are highly anti-correlated across all GCMs (for example $r = -0.91$ for SW-CRE over LLO) and is attributable to uncertainty in the regression intercept (effective radiative forcing) which can be strongly effected by internal variability in the first few years following the doubling of $CO_2$ concentrations, particularly on regional scales. To effectively distinguish physical fast adjustments from temperature dependent feedbacks on a regional scale requires a greater signal to noise ratio (such as a quadrupling of $CO_2$ or an initial condition ensemble) than used here and is work that I am currently undertaking with these GCMs.

For all subsequent analysis within this chapter, I choose to define the climate response as the change in the state of the climate system averaged over the final 20 year period available in the STEP experiment timeseries, normalised by the average GMST change over this period [Cess *et al.,* 1990] to avoid issues surrounding the physicality of partitioning climate response into forcing and feedback. This choice has the advantage of linking CRE responses directly to the changes in forcing governing the warming of the GCM at the end of the integration period.

Figure 6.8 shows normalised CRE changes across the local ensemble, decomposed by region. Regions are defined identically to chapter 5. Differences in CRE response between the GCMs can (as in chapter 5) be attributed to differences in the SW-CRE response over LLO, with additional smaller contributions from both SW-CRE and LW-CRE cooling of the climate system across other oceanic regions. The standard HadCM3 (black line) shows very close to zero SW-CRE response in both low and

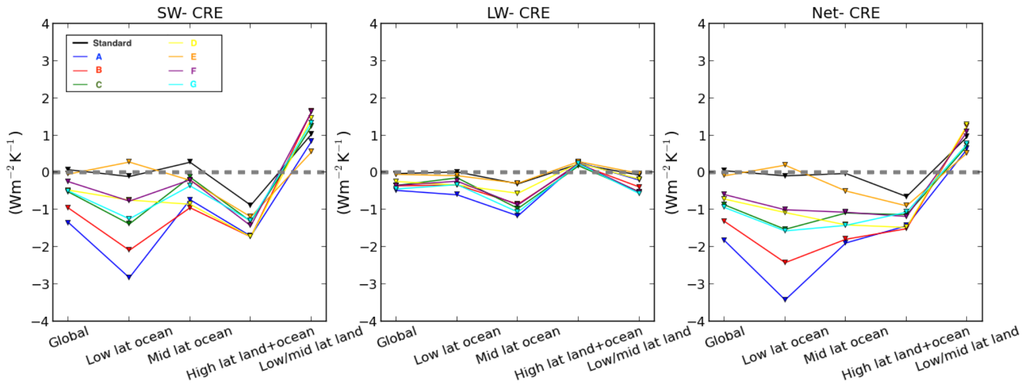


**Figure 6.8** – Regional cloud feedbacks. Panels show changes in components of CRE averaged over different geographical regions for the final 20 years for of STEP integration, normalised by the change in GMST. Regions are defined as: Low lat ocean (LLO - ocean; 30°N-30°S), Mid lat ocean (MLO - ocean; 55°N-30°N and 30°S-55°S), High lat land+ocean (HLLO - land and ocean; 90°N-55°N and 55°S-90°S) and Low/mid lat land (LML - land; 55°N-55°S).

mid-latitude oceans.

Focusing again on the global scale, relative humidity responses (not shown) are robustly simulated across the local ensemble with a clear rising of the tropopause in the tropics and a moistening of the upper troposphere in all GCMs. The similarity of this pattern of relative humidity change indicates robust features of circulation response across the GCMs that are insensitive to parameter perturbations, including an expansion of subsidence regimes associated with a broadening of the Hadley Cell. Changes in relative humidity are fractionally small compared to the mean state of the control climate and are generally co-located with changes in cloud fraction in the free troposphere.

Figure 6.9 shows global maps of CRE response for the local ensemble. Strong increases in SW-CRE cooling can be observed in all GCMs with low TCR over oceanic regions typical of shallow cumulus/stratocumulus cloud cover and large-scale subsidence. Figure 6.10 shows that these areas are strongly co-located with areas of low cloud growth in the warmed climate, suggesting that increases in the spatial coverage (as opposed to increases in cloud optical depth) at least partly drive the SW-CRE response in these regions. The spatial structure of the SW-CRE response is complex, with decreases in SW-CRE cooling simulated in the prevalent oceanic stratocumulus decks off the west coast of sub-tropical continents and adjacent to regions of increased cooling. Changes in LW-CRE are typically negative, but display a much wider variety of patterns between the members of the local ensemble, and are smaller in magnitude than the SW-CRE response in all perturbed GCMs.

Net-CRE response patterns have strong couplings to changes in surface air temperature in the tropics (figure 6.11). For a region in the Pacific, south of the equator, which experiences substantial growth in SW-CRE cooling in all GCMs of low TCR, local surface air temperatures actually reduce despite globally-mean warming. The magnitude of this response appears unconnected with (and is uncorrelated with) the control climate mean state of the surface temperature in this location (left hand side of figure 6.11).

### 6.3.1 Regime specific responses of tropical clouds

The importance of tropical cloud feedbacks in determining GCM sensitivity has been recognised in many previous studies (e.g. Bony *et al.* [2006] or Webb *et al.* [2006]), as have the importance of global SW cloud feedbacks [Vial *et al.,* 2013; Donohoe *et al.,* 2014]. In particular, several studies have argued that the greatest uncertainty in tropical cloud feedbacks can be traced to responses of low-level marine boundary layer clouds [Bony & Dufresne, 2005; Webb *et al.,* 2013; Vial *et al.,* 2013].

To further assess this complex pattern of CRE response, and the relative importance of different tropical cloud regimes, I adopt a stability indexing approach to isolate different regimes of tropical cloud behaviour. Many different stability binning indices have been used in the literature including, the vertical pressure velocity at 500hPa ($\omega_{500}$) [Bony & Dufresne, 2005], the Lower Tropospheric Stability (LTS - difference between the potential temperature at 700hPa and the surface) [Klein & Hartmann, 1993] and the Estimated Inversion Strength [Wood & Bretherton,

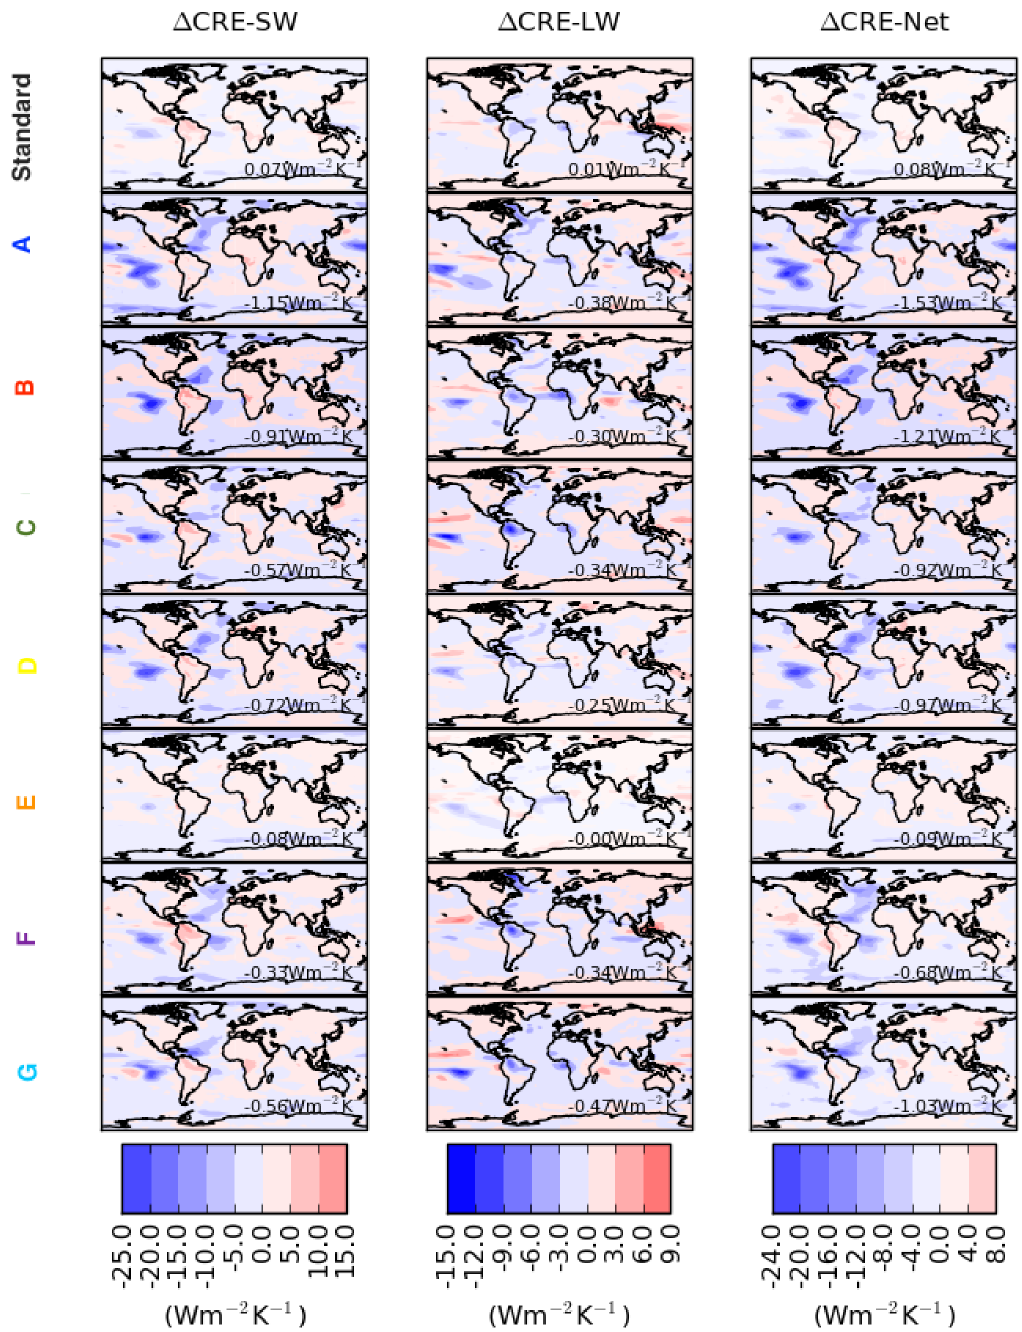


**Figure 6.9** – Changes in components of CRE normalised by the changes in GMST. The numbers given in each panel show the global average of the feedback.

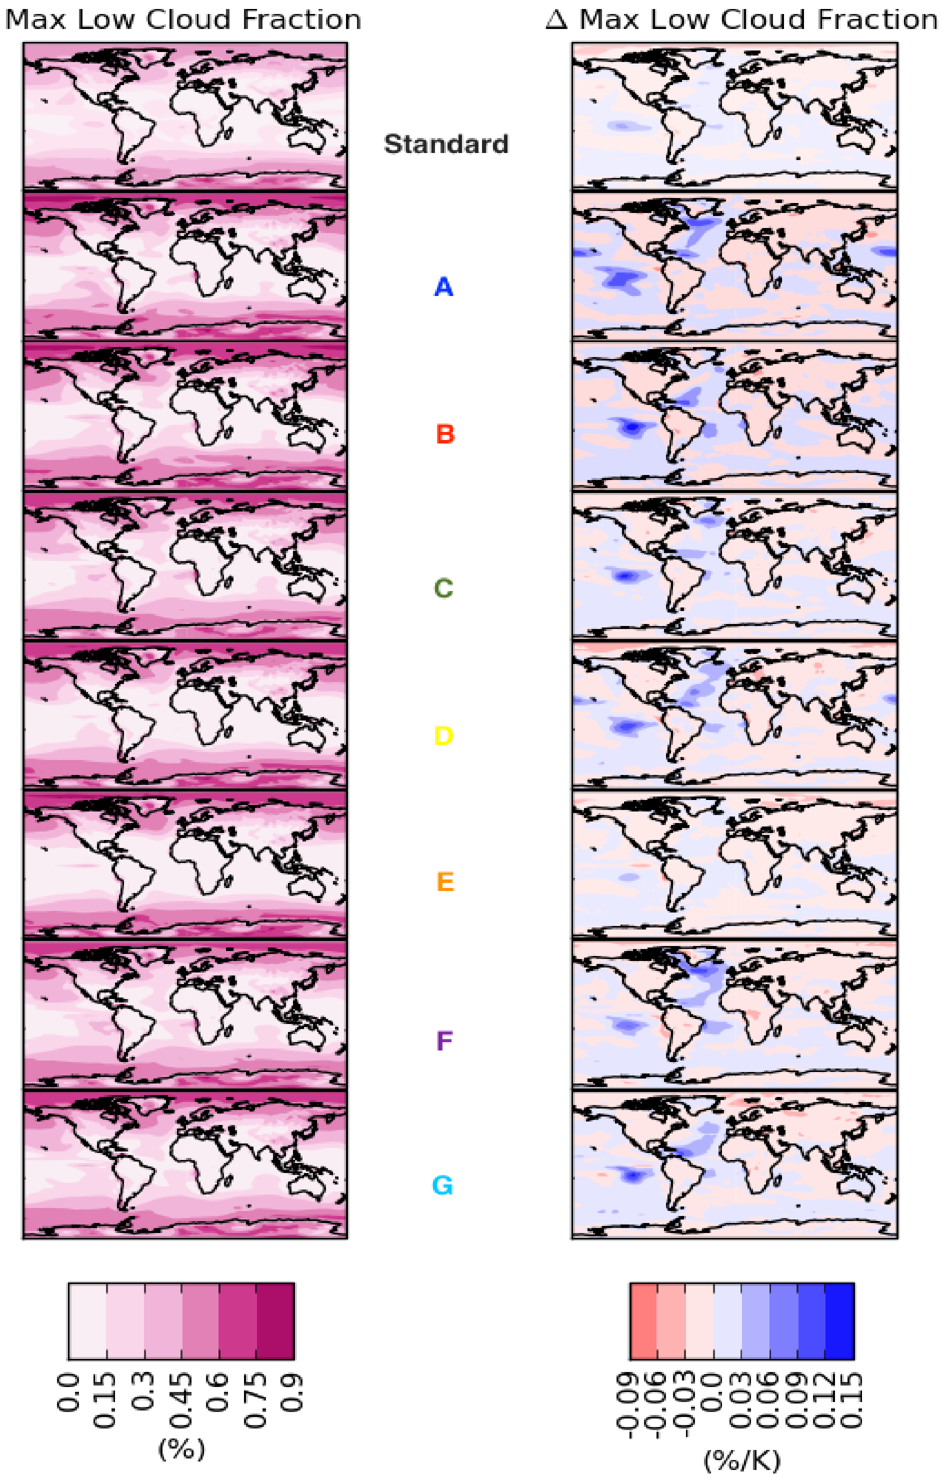


**Figure 6.10** – Control climate low cloud (>860hPa) fraction (left hand side) and the normalised change in low cloud fraction per degree of GMST warming (right hand side).

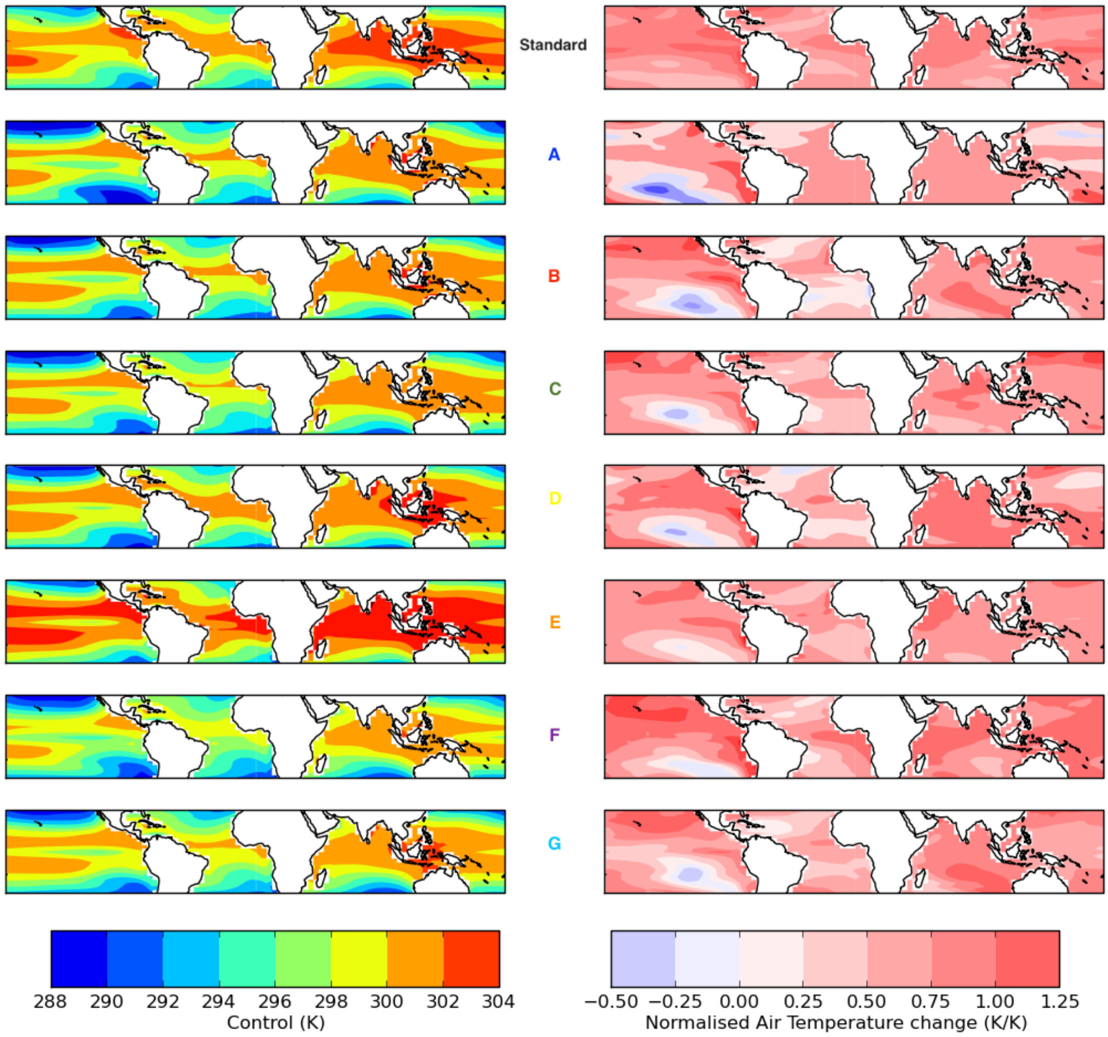


**Figure 6.11** – Left, the control climate surface air temperature and right, changes in surface air temperatures per degree of GMST warming under the doubled $CO_2$ experiment.

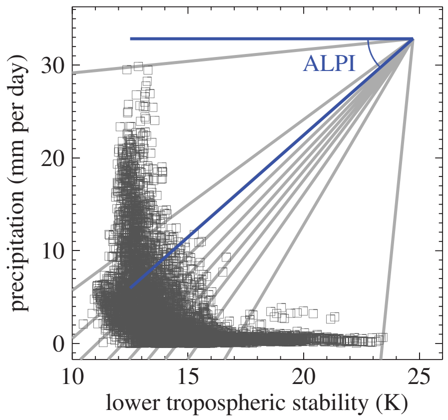


**Figure 6.12** – Figure 2 reproduced from Webb *et al.* [2015], showing the definition of the ALPI index.

2006]. Webb *et al.* [2015] developed a new stability binning index named the *angular LTS/precipitation index* (ALPI - see figure 6.12). This index recognises that a wide variety of precipitation is possible at very similar values of the LTS, and acts as a composite of information about the stability of the atmosphere from both the local LTS and precipitation fields. ALPI is evaluated on a monthly basis and is defined as,

$$ALPI = \arctan\left(\frac{P_{distance}}{LTS_{distance}}\right), \tag{6.3.1}$$

where $P_{distance}$ is defined as,

$$P_{distance} = \frac{P_{max} - P}{P_{max} - P_{min}} + 0.1, \tag{6.3.2}$$

and $LTS_{distance}$ is defined as,

$$LTS_{distance} = \frac{LTS_{max} - LTS}{LTS_{max} - LTS_{min}} + 0.1. \tag{6.3.3}$$

Due to its composite nature, ALPI is used as an index for tropical cloud regimes throughout the rest of this chapter. Results presented in this section are not especially sensitive to the choice of stability index with LTS and EIS binning giving broadly similar pictures of cloud responses (not shown).

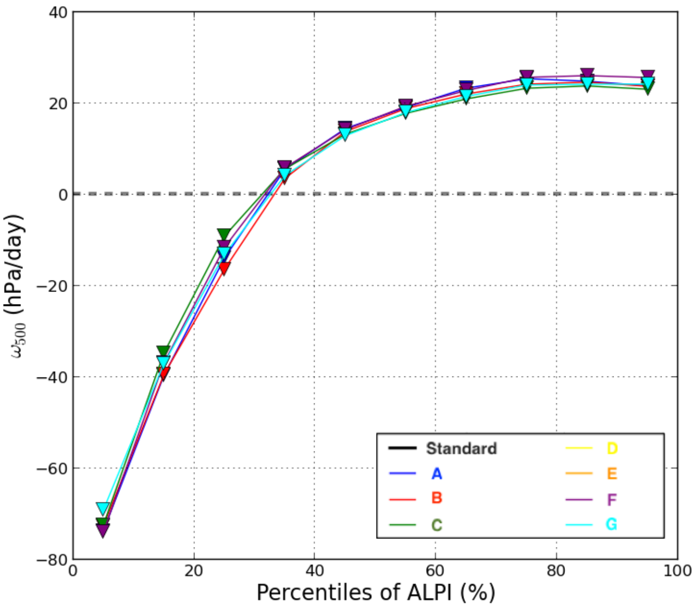


**Figure 6.13** – $\omega_{500}$ averaged over different ALPI deciles over LLO in the local ensemble GCMs.

To bin geophysical fields according to their ALPI simulated stability, ALPI is evaluated at every grid-box for each month of the simulation and then grouped into deciles of the ALPI index, where each bin is chosen such that it covers 10% of oceanic surface area between 30°N and 30°S. Binning is done separately for each month of the simulation and taking a weighted average therefore gives a timeseries for each decile of the ALPI distribution. The advantage of binning with equal areas in each bin is that bin sizes will automatically remain constant between control and STEP experiment climates, unlike when using fixed intervals of an index (Webb *et al.* [2013] and Webb *et al.* [2015] discuss this point in more detail). The connection with the large-scale vertical motion of the atmosphere is shown in figure 6.13, where it is clear that the upper end of the ALPI distribution (>50th percentile) correspond to the regimes of weak subsidence (approximately 20hPa/day) that are used in many studies of the low cloud response [Brient & Bony, 2012; Brient & Bony, 2013].

Figure 6.14 shows properties of ALPI deciles for local ensemble control climates. Low percentiles of the ALPI distribution correspond to areas of deep convection

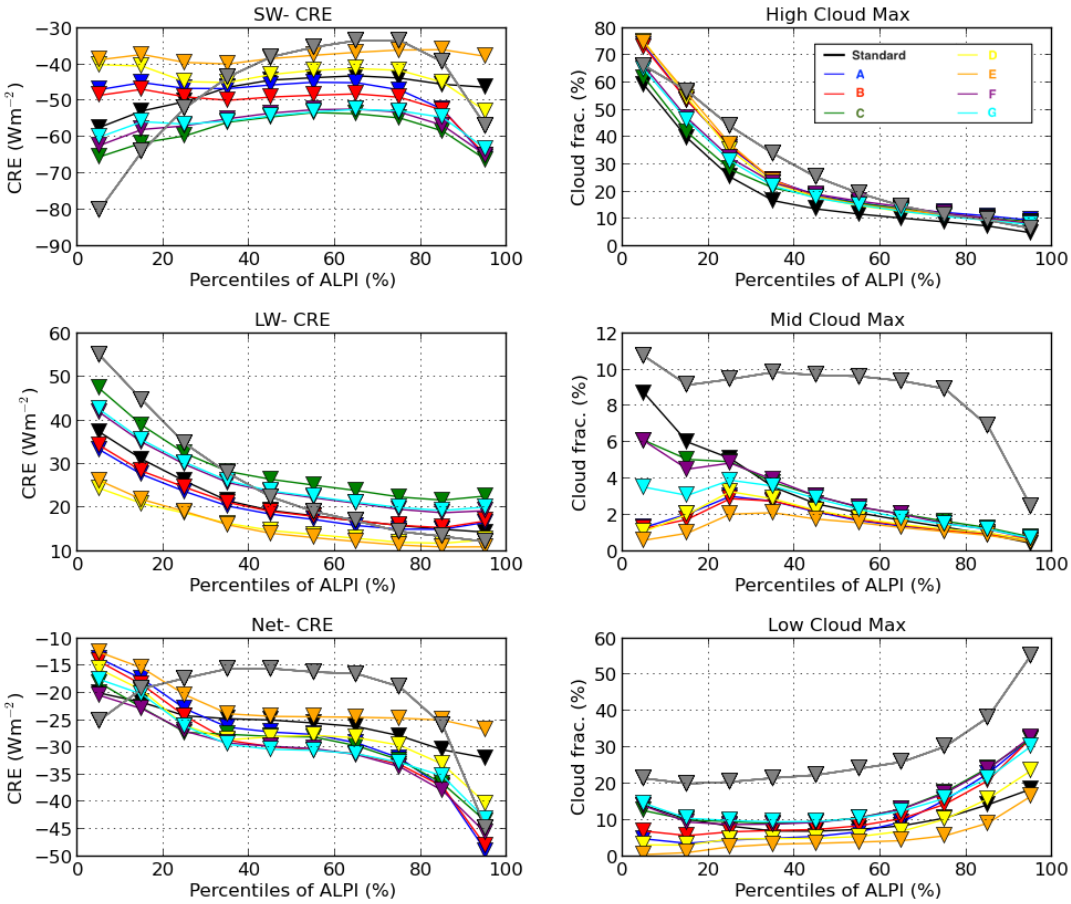


**Figure 6.14** – Cloud cover and radiative fields binned by ALPI deciles in control simulations for low-latitude oceans (30°N-30°S). The left hand column shows CRE components and the right hand shows: high cloud cover (max cloud fraction on pressure levels <440hPa; first row), mid cloud cover (between 440 and 680hPa; middle row) and low cloud cover (>680hPa; bottom row). ALPI deciles are plotted at the midpoint of each bin on the x-axis.

with high precipitation and are mainly located in the ITCZ. High percentiles of the ALPI distribution mainly correspond to subsidence regions that are generally populated with low level boundary layer clouds. Across the local ensemble, a substantial difference is observed between the standard HadCM3 (black line) and some of the lower TCR GCMs in their simulation of the high (stable) end of the ALPI distribution. In particular, these GCMs have more low clouds and increased SW-CRE cooling from these stable regions of the tropics. Additionally, reductions in low cloud cover relative to standard HadCM3 in the lowest decile (most unstable) of ALPI also decreases SW-CRE cooling in these deep-convection areas. High cloud fractions are simulated consistently across all GCMs, but mid cloud fractions are typically lower than in the standard HadCM3.

Figure 6.14 shows (grey) a combined observation/reanalysis measure of the present climate. The ALPI index is calculated using climatological temperature and pressure from the ERA-Interim reanalysis and climatological precipitation from the GPCP project, CREs are evaluated from the CERES observations with cloud fractions at different altitudes again taken from the ERA-Interim reanalysis. All these fields are re-gridded to the HadCM3 grid and evaluated over the common-subset of the periods used in section 6.2.

LLO high cloud fractions are simulated well by the members of the local ensemble, particularly in areas of subsidence (high ALPI). In general the local ensemble members improve the simulation of high clouds in areas of strong convection (low ALPI) relative to standard HadCM3. Substantial errors exist in mid-level cloudiness for all GCMs in the local ensemble. Over regions of low and mid ALPI all GCMs (including the standard HadCM3) have too little mid-level cloud compared to the reanalysis, whereas the standard HadCM3 is much better at capturing mid-level cloud fractions in regions of deep convection than the perturbed GCMs. A common feature of all GCMs in the local ensemble is insufficient low cloud coverage across all regimes of the LLO atmosphere. In general, low TCR GCMs improve this bias and manage to capture the non-monotonic nature of the SW-CRE at high ALPI

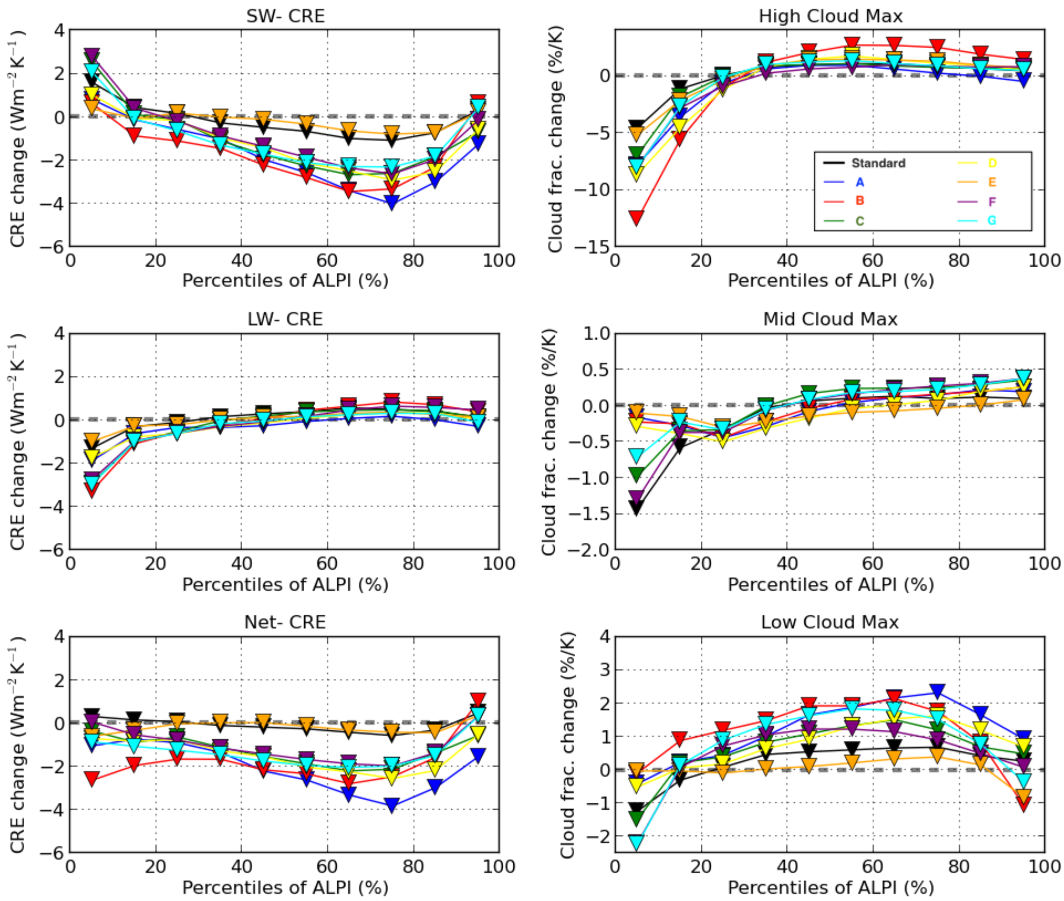


**Figure 6.15** – As for figure 6.14 but now expressed for changes in the same quantities.

(unlike standard HadCM3). LW-CRE is typically too small in all GCMs for deep-convective areas. CMIP5 GCMs, analysed in Webb *et al.* [2015], show similar biases in mid level cloudiness and smaller biases of both LW-CRE and SW-CRE in low ALPI/high convection regions. Simulation of cloud fractions at different altitudes are largely similar between these GCMs and the CMIP5 GCMs, again indicating that these GCMs do not have tropical cloud climatologies that would be considered complete outliers from the CMIP5 distribution and perform as well, if not better than many CMIP5 models for certain measures of tropical cloud simulation skill.

Figure 6.15 shows changes in the same quantities as in figure 6.14 at the end of the STEP experiment, normalised by GMST change. Low cloud fractions increase markedly with warming in many of the low climate response GCMs, over the mid-to-upper ALPI deciles. Mid-level cloud changes are generally small in the local ensemble aside from over the most unstable (low) ALPI decile. Tropical cloud radiative feedbacks are driven by increased SW-CRE cooling over mid-to-high ALPI deciles. Increases in low cloud are highly correlated ($r = 0.95$) with reductions in SW-CRE across ALPI deciles in the upper half of the range, suggesting that SW-CRE changes are driven, at least in part, by increasing average cloudiness as opposed to an increase in the optical density of low clouds. High cloud fractions decline across low ALPI declines in the perturbed GCMs and particularly over regions of deep convection (low ALPI), whereas high cloud fractions show a smaller reduction under warming in standard HadCM3. The picture given in figure 6.15 is consistent with an increase of SW-CRE cooling of the climate system driven by increases in low level cloudiness over the tropical oceans by the end of the STEP experiment.

## 6.4 Physical mechanisms of low climate response

The previous section described differing increases of low cloud in shallow cumulus/stratocumulus regions with warming, driven primarily by variations in SW-CRE feedbacks across the local ensemble. This section investigates potential physical mechanisms underlying this diversity of climate response.

The response of planetary boundary layer (PBL) clouds under a perturbation of $CO_2$ has traditionally been difficult to simulate consistently in GCMs due to the large range of dynamical and microphysical processes that underlie the representation of PBL clouds (see Bretherton [2015] for a review using high resolution climate models). Many of these small-scale processes, connected to convective and turbulent movements of moisture in the PBL, have to be parametrised at the gridscale in global GCMs. While there is no clear distinction between convective and turbulent

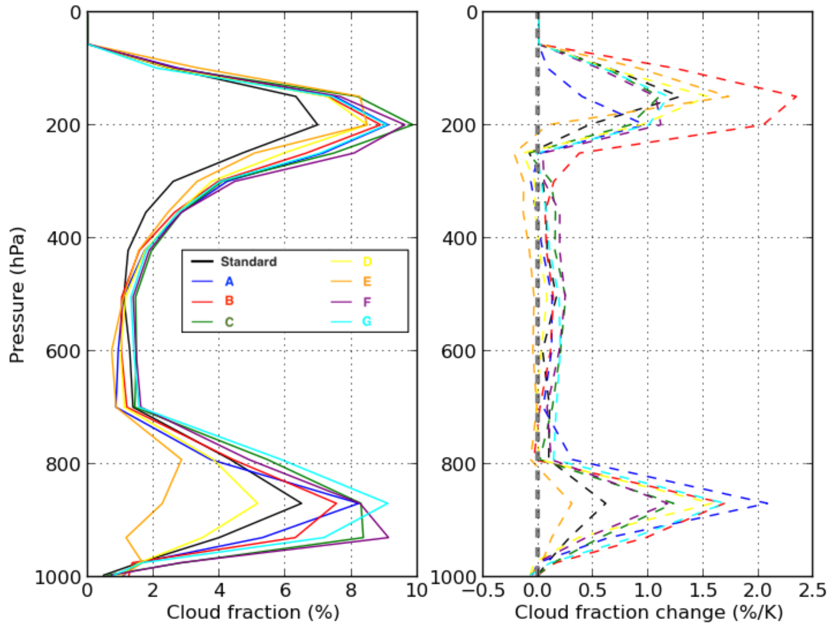


**Figure 6.16** – Cloud fraction on GCM levels (left hand panel) and changes per unit GMST change over the 50th-90th percentiles of the ALPI index (right hand panel).

motions, their effects are artificially separated in GCMs into contributions from the convection, and boundary layer turbulence parametrisation schemes.

Many competing hypotheses exist to explain how these clouds are simulated to respond to warming in GCMs. Wyant *et al.* [2009] argue for the increased humidity in the PBL of subsidence regimes driving increased clear-sky radiative cooling, destabilising the PBL and allowing the low cloud fraction to increase. However Webb & Lock [2013] suggest that the PBL specific humidity may decrease under warming, leading to a loss of low cloud cover atop the PBL.

Figure 6.16 shows control climate cloud fractions (left hand panel) and the changes in cloud fractions on GCM levels normalised by the GMST warming (right hand panel). Averages are taken over the 50th-90th percentiles of the ALPI distribution, where SW-CRE response to global warming is the most pronounced. Cloud fractions for low clouds, which sit atop the PBL, are typically greater in the perturbed GCMs than the standard HadCM3. The greatest diversity amongst the GCMs is in their representation of low and high cloud fractions, with very little diversity in

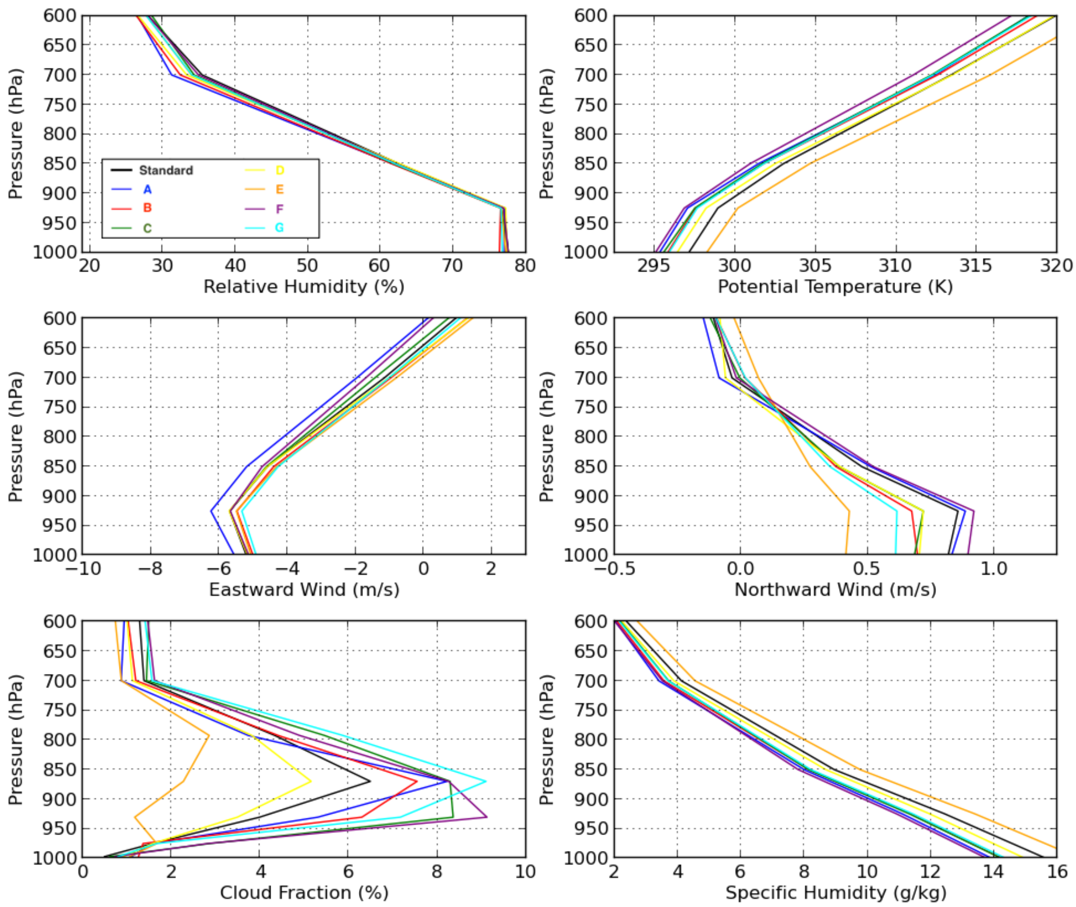


**Figure 6.17** – Annual mean vertical profiles averaged over the 50-90th percentile ALPI bin.

mid-level cloud fractions apparent in either control climate or climate response. In all GCMs a robust increase in the height of the high cloud maximum is observed, consistent with an approximately constant cloud top temperature for deep convection (due to the need to balance radiative cooling and convective heating in the tropics and the strong dependence of water vapour cooling on temperature Hartmann & Larson [2002]). Whilst high cloud changes do show differences between the perturbed GCMs, they are not significantly correlated ($r = -0.06$) with the GCM TCR (unlike the increase in low cloud fraction for which $r = -0.96$) and therefore I choose not to further analyse the effects of high cloud feedbacks in this chapter. The lack of mid-level cloud responses further supports low cloud responses as the prime determiner of SW-CRE feedbacks in the perturbed GCMs.

What might be underlying these differences in low cloud response? Figure 6.17 shows vertical profiles for several fields through the lower atmosphere averaged over the 50-90th percentile of the ALPI distribution. Specific humidity is derived from the relative humidity and temperature fields using a Magnus approximation for saturation vapour pressure [Alduchov & Eskridge, 1996]. The higher TCR GCMs (black and orange) are typically warmer at the surface than the lower response GCMs although no significant correlation exists across the local ensemble. Vertical relative humidity profiles are generally similar across the set of GCMs. As discussed previously, large differences exist in control climate low cloud fractions across the local ensemble. Low cloud cover has a similar vertical structure in all local ensemble GCMs with the maximum of cloud fraction typically located between 900-850hPa (one GCM has a maximum closer to 950 hPa and one above 800hPa). Boundary layer profiles of potential temperature and relative humidity are less uniform in the vertical than in CMIP5 GCMs due to the reduced vertical resolution of GCM layers near to the surface of the Earth in HadCM3.

Figure 6.18 shows the changes in vertical profiles through the atmosphere averaged over the 50-90th percentile of ALPI and normalised here by the *local* surface air temperature change over the same ALPI deciles only. The level of maximum cloud fraction increase is associated with an increase in local relative humidity at the same height in all members of the low sensitivity ensemble. An increase in local atmospheric stability (changes in vertical gradients of potential temperature) is seen in all GCMs but differences in magnitudes exist across the local ensemble. This robust increase in stability has been argued to suppress convective mixing (both moistening of the boundary layer through shallow convection from the surface, and drying via increased entrainment of dry free-troposphere air), leading to increased cloud fraction [Miller, 1997]. The members of the local ensemble show changes in the vertical gradient of specific humidity that are more similar than their changes in vertical potential temperature gradient, suggesting that humidity tendencies are less important for the simulation of low cloud growth than the temperature struc-

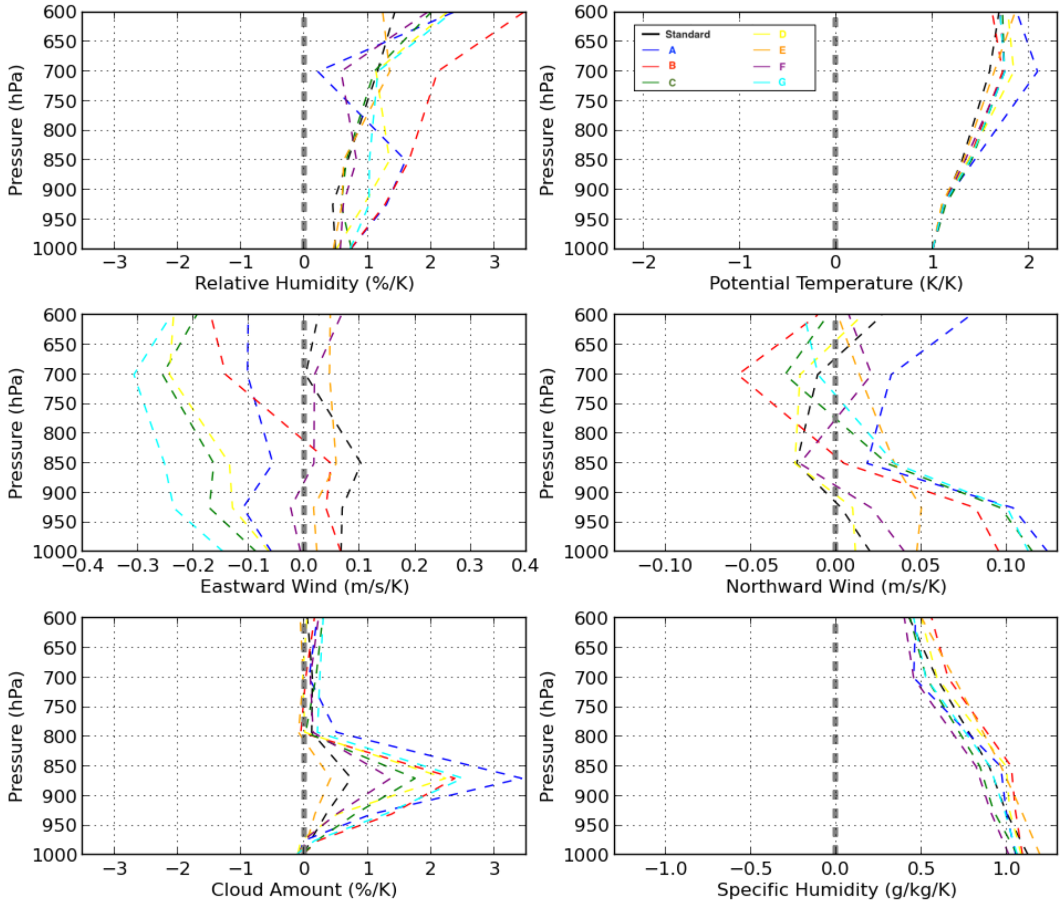


**Figure 6.18** – As for figure 6.17 but for changes in the same quantities normalised by the surface air temperature change over the same ALPI deciles only.

ture mediated effects in these GCMs. In all members of the ensemble, surface latent heat fluxes increase in response to warming (not shown), suggesting an increase in turbulence driven moistening of the boundary layer associated with increases in the specific humidity vertical gradients near the surface (larger specific humidity gradients lead to a greater moisture flux when two parcels of air are exchanged vertically under turbulent motions in the PBL). As increases in specific humidity gradients are similar across the ensemble, it therefore indicates that it is likely that the differing magnitudes of convective drying of the PBL, through entrainment of dry free troposphere air (mediated through differing changes in potential temperature gradients), is the prime determiner of the differing rates of low cloud growth, as opposed to turbulent driven entrainment of free-troposphere air through the capping boundary layer inversion. The perturbed GCMs all simulate a decrease in subsidence over these regions of the LLO, adding an additional large-scale resolved motion moistening tendency to those derived from the parametrisation schemes due to more vertically extensive (and therefore thicker) PBL clouds under weaker subsidence.

Additional influences on PBL clouds can arise from direct effect of $CO_2$ in the atmosphere, connected to changes in the net radiative cooling of the cloud tops. Gregory & Webb [2008] highlight the existence of substantial tropospheric fast adjustments in CRE across the CMIP3 MME. An increase in the concentration of $CO_2$ without an associated induced warming acts to reduce the net long-wave PBL cloud top radiative cooling that in turn could reduce vertical convective movement of water vapour by acting as a stabilising effect on the PBL. Kamae & Watanabe [2012] find tropospheric fast adjustments that both increase and decrease (although predominately decrease) low cloud cover in stable regimes of the tropics in the CMIP5 ensemble when $CO_2$ concentrations are increased and warming is prevented. Fast tropospheric adjustments may play an important role in determining the changes in low clouds observed at the end of the doubled $CO_2$ integrations for the local ensemble GCMs. However, as discussed previously, higher signal to noise is required to clearly identify fast adjustment effects on cloud from their response to surface warming and will be an area of further investigation with these GCMs.

Sherwood *et al.* [2014b] hypothesise that across GCMs rates of convective mixing between the boundary layer and free-troposphere in the present climate are correlated with the strength of the response of this process to climate change. They suggest that GCMs which dehydrate the boundary layer more in the present climate increase (or decrease less) their boundary layer dehydration more in a warmed climate. They define indices of convective and large-scale mixing between the mid and lower troposphere for regions of tropical deep convection (see Sherwood *et al.* [2014b]) and show correlations between these indices and the ECS of several CMIP5 GCMs. Table 6.2 shows the same indices evaluated for members of the local ensemble for which all the required fields were available.

**Table 6.2** – Sherwood *et al.* [2014b] indices for the local ensemble GCMs with all required fields available.

| GCM | S | D | S+D |
|---|---|---|---|
| A | 0.35 | 0.18 | 0.53 |
| B | 0.35 | 0.15 | 0.50 |
| C | 0.38 | 0.18 | 0.56 |
| F | 0.39 | 0.19 | 0.58 |
| G | 0.36 | 0.16 | 0.52 |

$S$ is a measure of the small-scale (convective) moisture mixing between the middle and low troposphere defined as,

$$S = (\Delta R_{700\text{hPa}-850\text{hPa}}/100\% - \Delta T_{700\text{hPa}-850\text{hPa}}/9\text{K})/2, \qquad (6.4.1)$$

where $\Delta R_{700\text{hPa}-850\text{hPa}}$ and $\Delta T_{700\text{hPa}-850\text{hPa}}$ are relative humidity and temperature differences between 700hPa and 850hPa respectively. Small relative humidity differences are associated with large mixing between the lower and mid troposphere and with more negative temperature differences (if air was directly lifted from the PBL and detrained directly from the plume the air would be moister and cooler than air that had risen high in deeper convection, warmed and dehumidified via

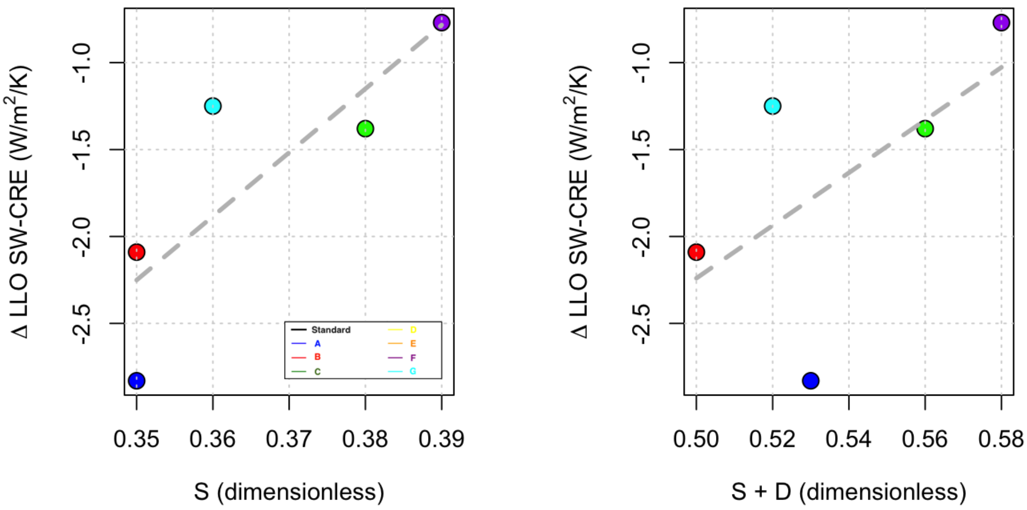


**Figure 6.19** – Sherwood *et al.* [2014b] S index (x-axis - left hand panel) and S + D index (x-axis - right hand panel) against GMST change normalised LLO SW-CRE response. The dashed grey lines show a ordinary least squares linear fit in each panel.

condensation/precipitation, and then subsequently descended back to just above the PBL). $S$ is calculated over all LLO grid-boxes which are in the upper quartile of $\omega_{500}$, corresponding to deep convective regions despite the main differences in cloud feedbacks arising in the subsidence regions, hypothesising that the index captures consistent tropical-wide strengths of shallow convection.

$D$ captures the large-scale resolved contribution to the mixing between lower and mid troposphere and can be defined as,

$$D = \frac{\Delta H(K)H(-\omega_1)}{-\omega_2 H(-\omega_2)}, \tag{6.4.2}$$

where $\omega_1$ is the average of $\omega$ at 700hPa and 850hPa, $\omega_2$ is the average of $\omega$ at 600hPa, 500hPa and 400hPa, $K = \omega_2 - \omega_1$ and H is a step function. $D$ is indicative of where moisture is being transported upward (and therefore also outward) by the large-scale circulation, as the net flow of air is upwards when $\Delta > 0$ and $\omega_1 < 0$. A summation of $S$ and $D$ therefore gives an index approximating the total mixing between the lower and mid-troposphere.

Figure 6.19 shows a positive correlation between the index $S$, or the combination

$S+D$, and the LLO SW-CRE feedbacks. GCMs that have stronger mixing between the low and mid troposphere, have the least negative SW-CRE feedbacks across the LLOs, consistent with the hypothesis of Sherwood *et al.* [2014b]. The correlation with SW-CRE is stronger for $S$ ($r = 0.82$) than $S + D$ ($r = 0.60$), indicating that convective mixing between the lower and mid troposphere is more important to the SW-CRE feedbacks than large-scale mixing. Sherwood *et al.* [2014a] hypothesise that this mixing is expected to increase under warming (or decrease less for particularly low present day mixing rates) in proportion to its present day strength, giving rise to the correlation with SW-CRE. Values of $S$ within the perturbed GCMs lie at low end of the range of the same index across the CMIP5 GCMs examined by Sherwood *et al.* [2014b], and would be associated with ECS of roughly 2K under the linear fit between $S$ and ECS in the CMIP5 GCMs. It must be noted that if this mechanism can be convincingly shown to be a physically sound emergent constraint on ECS and cloud feedbacks, the observationally-constrained value for S of $\approx 0.4$ and $\approx 0.8$ for S + D, could be used to rule out the low climate response members of the local ensemble. Substantial scatter exists about the linear relationships shown in figure 6.19, indicating that this mechanism cannot, by itself, explain all of the SW-CRE changes simulated in the local ensemble GCMs.

## 6.5 Conclusions

In this chapter I have analysed a representative set of GCMs from the predicted low sensitivity ensemble of chapter 5. I analysed the control climates of these GCMs and found a diverse picture of how biases from observations differ in the perturbed versions of the GCM from the standard HadCM3. Some biases are improved in the perturbed GCMs, such as the coverage of low cloud over tropical oceans, SW-CRE over the most stable parts of the tropical oceans and the RMSE of the precipitation climatology. However, low TCR perturbed GCMs show, in general, lower skill compared to the standard HadCM3 and in particular for climatological representations of SW-CREs.

An important question to consider is whether the low climate response GCMs highlight any particular biases, the alleviation or extenuation of which may be connected with low climate response. All the low response GCMs show a reduction of biases in low cloud cover over the low-latitude oceans. The lack of low cloud over these regions is a common bias of GCMs. As the low climate response in the perturbed GCMs is strongly driven by negative feedbacks in low-latitude oceanic cloud, this raises an interesting question as to the role of the tropical low cloud cover bias in determining the GCM climate sensitivity. This would be an interesting and valuable topic of study with perturbed versions of other GCMs.

Turning to features of the climate response, I have shown that low sensitivity GCMs are associated with very strong increases in SW-CRE cooling over subsidence areas of low-latitude oceans. An increase in the low cloud cover is seen with warming in these regions under doubling of $CO_2$. Low cloud feedbacks have repeatedly been found to be the greatest source of uncertainty in ECS [Webb *et al.,* 2013], and therefore the identification of the same feedback as important for determining the climate response in the perturbed GCMs shows that the low climate responses in these GCMs are not achieved through a new mechanism unknown in other GCMs.

Understanding the physical mechanisms driving the simulated increase in low clouds is a challenging task, as with other GCMs that have been previously studied [Webb & Lock, 2013]. This is particularly a challenge in a GCM with low vertical resolution in the PBL, such as HadCM3, as much of the important physics needs to be parametrised on relatively few vertical layers. The mechanisms of changes in mid-to-low tropospheric mixing under warming proportional to the strength of this mechanism in the control climate, hypothesised in Sherwood *et al.* [2014b], seems to explain some of the variance in climate response across the local ensemble. Whilst this hypothesis seems consistent with an analysis of vertical atmospheric profiles, the causality of the physical responses underlying these changes can be difficult to deduce. The computation of tendencies on vertical levels in the atmosphere, decomposed by parametrisation scheme, for different thermodynamic variables, would offer valuable further insights into convection and turbulence mediated effects in driving this response. Similarly, idealised experiments in which features of the climate response are artificially suppressed, such as fixed lower tropospheric stability change, could separate the effect of different mechanisms on the simulated changes in the local ensemble GCMs. A final useful extension would be an enhancement of the climate change signal-to-noise ratio by conducting quadrupling of pre-industrial $CO_2$ experiments in order to effectively distinguish physical tropospheric fast adjustments from temperature-mediated feedbacks. I am currently beginning extension work along several of these lines with the GCM parameter sets analysed in this chapter.

CHAPTER 7

# Simple carbon-cycle representations and adaptive mitigation policies under climate response uncertainty

This thesis has thus far focused on physical uncertainties in the climate response to radiative forcing. Chapter 1 provided a simple example of how physical climate response uncertainty can impact the expected cost of climate change induced economic damages (section 1.3). Previous chapters of this thesis have examined uncertainties in the climate response using different methods and concluded, based on the experiments undertaken, that insufficient evidence exists to further narrow existing uncertainty estimates of the climate response to $CO_2$. Dealing with endemic and substantial uncertainty in the physical climate response is therefore likely to remain a significant obstacle for attempts to construct mitigation policies aiming to limit warming to beneath a specific threshold.

This chapter proposes a methodology, in which global mitigation policies could adapt to the emerging climate response, to achieve a pre-defined warming goal. Section 7.1 describes the coupling of a simple global carbon-cycle model to the impulse-response simple climate model used in chapter 4 in order to construct an empirical climate-carbon-cycle model, named FAIR, for assessing uncertainty in the response of the climate system to $CO_2$ emissions. This model, by construction, captures the relevant physics of the land and ocean carbon-cycle response to anthropogenic emissions. Section 7.2 then describes and constructs an adaptive mitigation policy using the FAIR model to examine such a policy framing under a broad range of climate response uncertainty. I then discuss the economic and policy implications of this policy framing in section 7.3, before summarising the conclusions of this chapter (section 7.4).

The development and validation of the simple climate-carbon-cycle model developed in this chapter has been submitted to *Atmospheric Chemistry and Physics*. The adaptive mitigation policy proposal in this chapter has been submitted to *Environmental Research Letters*.

## 7.1 A modified impulse-response representation of the global response to $CO_2$ emissions

Physical climate science has historically focused on the response of the climate system to radiative forcing, by prescribing concentrations of $CO_2$ in the atmosphere. It is only with the most recent generation of ESMs that prescribed emissions runs with complex models have become widely available. Analysis of mitigation policy requires the additional step of connecting emissions of greenhouse gases to changes in atmospheric concentrations. Therefore, a model of carbon-cycle processes (that govern the evolution of the atmospheric concentration of $CO_2$), as well as the thermal response to radiative forcing, is required.

Simplified representations of the coupled climate-carbon-cycle system take many forms [Joos *et al.,* 1996; Meinshausen *et al.,* 2011; Huntingford *et al.,* 2009; Hof *et al.,* 2012]. A key test for simplified Earth system models is whether they correctly capture the physics of the co-evolution of atmospheric $CO_2$ concentrations and GMST under idealised settings. Following a $CO_2$ pulse emission of 100GtC in present-day climate conditions, ESMs (and Earth System Models of Intermediate Complexity - EMICs) display a rapid-draw down of $CO_2$ with the concentration anomaly due to the pulse reducing by approximately 40% from peak after 20 years and by 60% after 100 years, followed by a much slower decay of concentrations leaving approximately 25% of the peak concentration anomaly remaining in the atmosphere after 1000 years [Joos *et al.,* 2013]. This longevity of fossil carbon concentration anomalies, when combined with the long timescale thermal adjustment of the climate system [Held *et al.,* 2010], leads to a GMST rapidly increasing to a constant value (over approximately a decade) following a pulse emission of $CO_2$ [Joos *et al.,* 2013]. Warming does not noticeably decrease from this value over the following several hundred years, indicating that, short of artificial $CO_2$ removal (CDR) or other geoengineering methods, $CO_2$-induced warming is essentially permanent on human-relevant timescales.

### The IPCC AR5 Impulse-Response (AR5-IR) model

IPCC-AR5 proposed an idealised simple climate-carbon-cycle model for metric calculations, incorporating the "2-box" model of the temperature response to radiative forcing from chapter 4 with a "4-time-constant" impulse-response model of the $CO_2$ response to emissions [Myhre, 2013a]. This model represents the evolution of atmospheric $CO_2$ by partitioning emissions of anthropogenic $CO_2$ between four different reservoirs (all of which are empty in pre-industrial equilibrium) of atmospheric carbon anomaly ($R_i$) that each decay with a fixed time constant. Four carbon pools are determined to be sufficient to empirically represent the response of atmospheric $CO_2$ concentration anomalies following a pulse emission of 100GtC, above a specified background concentration of 389ppm, over the 1000 years following the pulse [Joos *et al.,* 2013]. These carbon pools do not directly correspond to individual physical processes and instead represent the combined effect of several carbon-cycle mechanisms, however, processes that are guiding analogues to the timescale of the pool decays are summarised in table 7.1. The impulse-response function for a unit emission at time $t = 0$ is given as,

$$\frac{\mathrm{d}R_i}{\mathrm{d}t} = a_i E - \frac{R_i}{\tau_i} \quad ; \quad i = 1 - 4 \tag{7.1.1}$$

where $E$ are annual $CO_2$ emissions, in units of ppm/year (1 ppm = 2.12GtC), $a_i$ is the fraction of carbon emissions entering each reservoir and $\tau_i$ the decay time constant for that pool. Atmospheric $CO_2$ concentrations are given by $C = C_0 + \sum_i R_i$ , and radiative forcing by:

$$F = \frac{F_{2\mathrm{x}}}{\ln(2)} \ln\left(\frac{C}{C_0}\right) + F_{\mathrm{ext}} \quad , \tag{7.1.2}$$

where $C_0$ is the pre-industrial $CO_2$ concentration, $F_{2\mathrm{X}}$ the forcing due to $CO_2$ doubling, and $F_{\mathrm{ext}}$ the non-$CO_2$ forcing. GMST anomalies are computed thus:

$$\frac{\mathrm{d}T_j}{\mathrm{d}t} = \frac{c_j F - T_j}{d_j} \quad ; \quad T = \sum_j T_j \quad ; \quad j = 1, 2 \tag{7.1.3}$$

with coefficients $a_i$, $d_j$ and $\tau_i$ as given in AR5 Chapter 8, tables 8.SM.9 and 8.SM.10 [Myhre, 2013a], except for $\tau_0$ which is here given the finite value and not set to infinity (all results presented in this chapter are insensitive to this choice). $c_j$ are set to give ECS=2.75K and TCR=1.6K (corresponding to $c_1 = 0.46\mathrm{KW}^{-1}\mathrm{m}^2$ and $c_2 = 0.27\mathrm{KW}^{-1}\mathrm{m}^2$ - see chapter 4), indicative of a typical mid-range climate response to radiative forcing in ESMs [Flato, 2013].

I use two versions of the AR5-IR model in this chapter, one calibrated to the present-day (AR5-IR) and one calibrated to the pre-industrial climate response to a pulse emission (PI-IR) respectively. The AR5-IR model is used for the calculation of absolute Global Temperature Potentials (aGTPs) in IPCC-AR5 and has carbon-cycle coefficients that best represent the evolution of a 100GtC pulse emission under approximately present-day conditions. The PI-IR model uses an alternative set of

| Parameter | Value - AR5-IR | Value - PI-IR | Value - FAIR | Processes |
|---|---|---|---|---|
| $a_0$ | 0.2173 | 0.1266 | 0.2173 | Geological re-absorption |
| $a_1$ | 0.2240 | 0.2607 | 0.2240 | Deep ocean invasion / equilibration |
| $a_2$ | 0.2824 | 0.2909 | 0.2824 | Biospheric uptake / ocean thermocline invasion |
| $a_3$ | 0.2763 | 0.3218 | 0.2763 | Rapid biospheric uptake / ocean mixed-layer invasion |
| $\tau_0$ (yr) | $1\text{x}10^6$ | $1\text{x}10^6$ | $1\text{x}10^6$ | Geological re-absorption |
| $\tau_1$ (yr) | 394.4 | 302.8 | 394.4 | Deep ocean invasion/equilibration |
| $\tau_2$ (yr) | 36.54 | 31.61 | 36.54 | Biospheric uptake / ocean thermocline invasion |
| $\tau_3$ (yr) | 4.304 | 4.240 | 4.304 | Rapid biospheric uptake / ocean mixed-layer invasion |
| $c_1$ ($\text{KW}^{-1}\text{m}^2$) | 0.46 | 0.46 | 0.46 | Thermal adjustment of upper ocean |
| $c_2$ ($\text{KW}^{-1}\text{m}^2$) | 0.27 | 0.27 | 0.27 | Thermal equilibration of deep ocean |
| $d_1$ (yr) | 8.4 | 8.4 | 8.4 | Thermal adjustment of upper ocean |
| $d_2$ (yr) | 409.5 | 409.5 | 409.5 | Thermal equilibration of deep ocean |
| $r_0$ (yr) | - | - | 35 | Pre-industrial $\text{iIRF}_{100}$ |
| $r_C$ (yr/GtC) | - | - | 0.02 | Increase in $\text{iIRF}_{100}$ with cumulative carbon uptake |
| $r_T$ (yr/K) | - | - | 4.5 | Increase in $\text{iIRF}_{100}$ with warming |

**Table 7.1** – Default parameter values for the simple impulse-response climate-carbon-cycle models used in this chapter. Note that, for consistency with [Myhre, 2013a], the ordering of indices is fast-slow for the thermal response and slow-fast for the carbon cycle.

coefficients that are selected to represent the evolution of a 100GtC pulse emission in pre-industrial conditions for the multi-model mean of the ensemble of ESMs and EMICs summarised in Table S2 of Joos *et al.* [2013] (see table 7.1 for parameter values).

### A Finite Amplitude Impulse Response (FAIR) model

In the AR5-IR model the carbon-cycle constants are not affected by rising temperature or $CO_2$ accumulation and hence only represent the specific response to a particular perturbation scenario. In more comprehensive models, ocean uptake efficiency declines with accumulated $CO_2$ in ocean sinks [Revelle & Suess, 1957] and uptake of carbon into both terrestrial and marine sinks are reduced by warming [Friedlingstein *et al.,* 2006].

In an attempt to capture some of these dynamics within the simple impulse-response model structure, I here attempt a minimal modification of the AR5-IR model to allow it to mimic the behaviour of ESMs in response to finite-amplitude $CO_2$ injections, which I will call a Finite Amplitude Impulse-Response (FAIR) model. To introduce a state-dependent carbon uptake as simply as possible, I apply a single scaling factor $\alpha$ to all four of the time-constants in the carbon-cycle of the AR5-IR model, such that the $CO_2$ concentration anomalies in the 4 "carbon reservoirs" are updated thus:

$$\frac{\mathrm{d}R_i}{\mathrm{d}t} = a_i E - \frac{R_i}{\alpha\tau_i} \quad ; \quad i = 1, 4 \tag{7.1.4}$$

To identify a suitable state-dependence, I focus on parametrising variations in the 100-year integrated impulse response function, $\mathrm{iIRF}_{100}$. A focus on the integrated impulse response (average airborne fraction over a period of time, multiplied by the length of time period), as opposed to the airborne fraction at a particular point in time, is more closely related to the impact of emissions on the global energy budget, and also to other metrics such as Global Temperature Potential (GTP) [Joos *et al.,* 2013]. With other coefficients fixed, $\mathrm{iIRF}_{100}$ is a monotonic (but non-linear) function of $\alpha$:

$$\mathrm{iIRF}_{100} = \sum_i \alpha a_i \tau_i \left[1 - \exp\left(\frac{-100}{\alpha\tau_i}\right)\right]. \tag{7.1.5}$$

Following other simplified carbon-cycle models [Meinshausen *et al.,* 2011; Glotter *et al.,* 2014], I assume $\mathrm{iIRF}_{100}$ is a function of accumulated perturbation carbon stock in the land and ocean (equivalent to the amount of emitted carbon that no longer resides in the atmosphere), $C_{\mathrm{acc}} = \sum_t E - (C - C_0)$, and of GMST anomaly from pre-industrial conditions, $T$. A simple linear relationship appears to give an adequate approximation to the behaviour of ESMs and EMICs (see section 7.1.1):

$$\mathrm{iIRF}_{100} = r_0 + r_C C_{\mathrm{acc}} + r_T T. \tag{7.1.6}$$

Whilst, the FAIR formulation does not separate carbon uptake into components from the land and ocean carbon cycles explicitly, the modelled temperature-induced changes in carbon uptake are primarily physically associated with responses of the land carbon-cycle [Arora *et al.,* 2013] as oceanic carbon uptake only displays a minor dependence on temperature via the solubility of $CO_2$ [Glotter *et al.,* 2014]. In contrast, direct $CO_2$ uptake and $CO_2$ concentration driven feedbacks are likely to be important both for the land and oceanic carbon-cycles.

At each time-step I first compute the required $\mathrm{iIRF}_{100}$ using $C_{\mathrm{acc}}$ and $T$ from the previous time-step (equation 7.1.6). I then numerical solve equation 7.1.5 to find the compatible value for $\alpha$, which is then in turn used to update the carbon pool concentration anomalies (equation 7.1.4). The total radiative forcing is then computed with equation 7.1.2, before changes in global mean temperature are computed with equation 7.1.3.

Values of $\mathrm{iIRF}_{100}$ larger than 100 years correspond to a net carbon source in response to a perturbation, and, as perturbations to the carbon stock in the atmosphere would grow indefinitely, makes the model unstable. In this regime there is no solution for $\alpha$, so I set $\mathrm{iIRF}_{100}$ to a maximum value of 95 years (chosen to approximately replicate the centennial timescale decay of concentrations in the UVic model following cessation of emissions in the A2+ scenario [Glotter *et al.,* 2014]), corresponding to $\alpha$=65.4. This physically corresponds to a near-absence of carbon sinks in the Earth system following a very large injection, with very slow rates of decay of atmospheric concentrations. The use of $\mathrm{iIRF}_{100}$ as a central measure of carbon-cycle response represents an inherent value choice about the timescale of climate-carbon-cycle system behaviour to be captured by the simple model. A time horizon of 100 years captures important aspects of the climate response to a pulse emission of $CO_2$ relevant over typical economic discounting timescales, whereas a longer time horizon could instead be used to prioritise the millennial scale response which may be more appropriate for some applications.

### 7.1.1 Testing the behaviour of the FAIR model

As a test of the behaviour of the FAIR model, I here conduct a set of experiments under idealised $CO_2$ emissions-driven and concentration-driven scenarios that have also been conducted in the literature for ESMs and EMICs. Throughout this section, I contrast the performance of the FAIR model, described in the previous section, to the AR5-IR model, PI-IR model, the MAGICC model [Meinshausen *et al.,* 2011] and the BEAM model [Glotter *et al.,* 2014], a simple carbon-cycle model that explicitly represents the physical effects of oceanic carbon uptake on ocean carbonate chemistry [Williams & Follows, 2011]. I use a version of the BEAM model with no temperature dependence of model parameters, which has been shown to be small in Glotter *et al.* [2014].

Gregory *et al.* [2009] conducted a set of experiments with two different ESMs under specified $CO_2$ concentration profiles (the ESMs are run with a prescribed concentration pathway as opposed to a prescribed emissions pathway), followed up with a broader range of models by Arora *et al.* [2013]. Concentrations were increased at 0.5%$yr^{-1}$, 1%$yr^{-1}$ and 2%$yr^{-1}$ respectively and consistent emissions (calculated as the residual of the carbon budget required to maintain the concentration timeseries after land and ocean carbon fluxes have been computed) were derived for different configurations of the ESMs: a "biogeochemically-coupled" experiment, where the carbon-cycle only responds to the direct effect of increasing $CO_2$ concentrations (as $CO_2$-induced warming is suppressed), a "radiatively-coupled" experiment in which the climate and carbon-cycle experiences $CO_2$-induced warming but the carbon-cycle is restricted to evolve as if the atmospheric concentration remains at the pre-industrial value, and a "fully-coupled" experiment in which the climate and carbon-cycle are allowed to respond to both warming and $CO_2$ concentration changes. Readers should consult Gregory *et al.* [2009] for further details on the above experiments.

Within the impulse-response modelling framework, I recreate the "biogeochemically-coupled" experiment by setting $r_T$ =0. I approximate the "radiatively-coupled" experiment by evaluating the difference between the "fully-coupled" and "biogeochemically-coupled" experiments. Although Gregory *et al.* [2009] found that the relationship between the experiments was not a simple linear summation at high $CO_2$ concentrations, a linear approximation serves as an adequate approximation for my purposes here, since the objective is the correct representation of aggregate feedbacks rather than a breakdown into specific contributions.

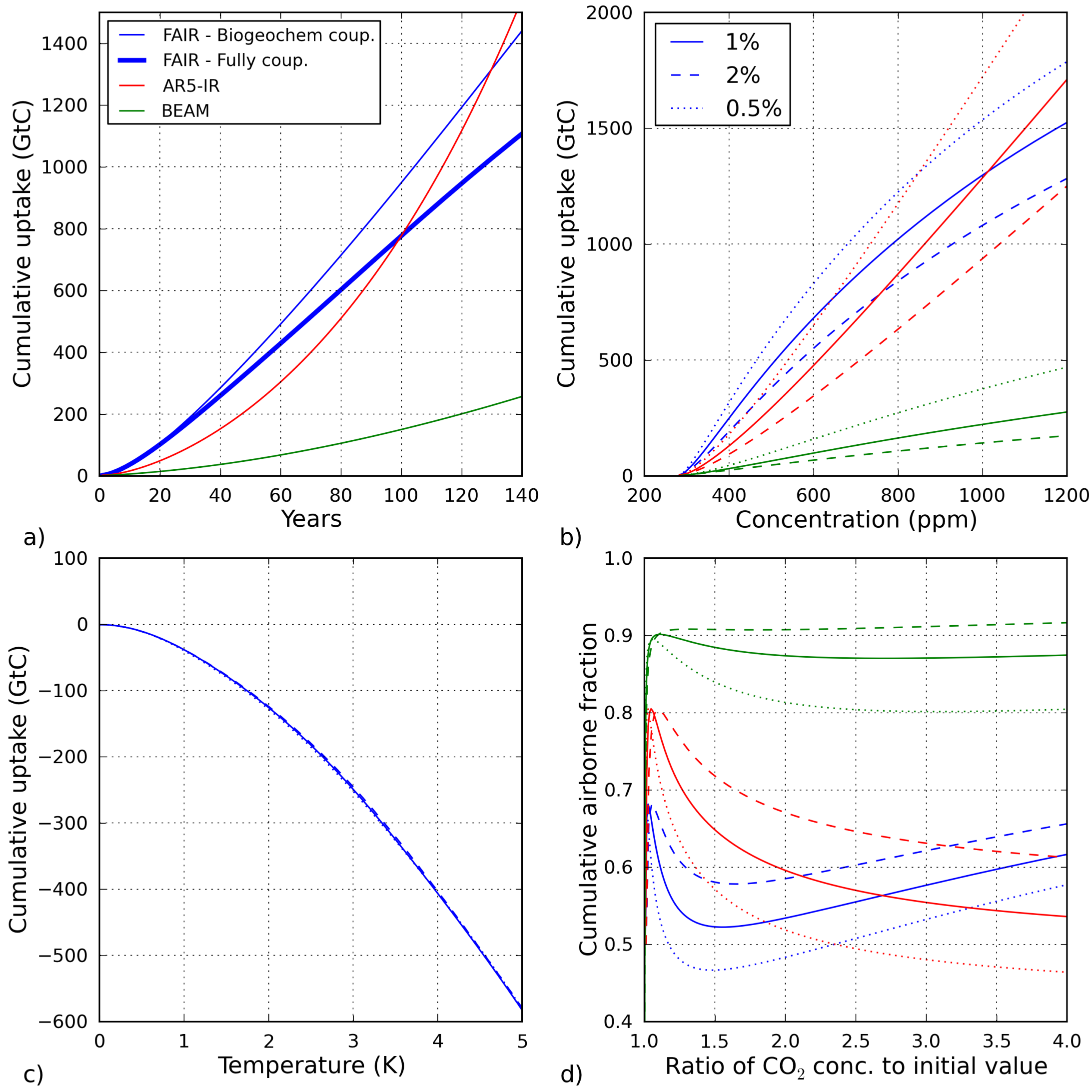


**Figure 7.1** – Response to idealised concentration increase experiments from Gregory *et al.* [2009] for the FAIR (blue), IPCC-AR5 (red) and BEAM (green) models. Panel a) shows the cumulative total carbon uptake over time in the "fully coupled" 1%$\text{yr}^{-1}$ concentration increase scenario. Panel b) shows the evolution of cumulative total carbon uptake as a function of atmospheric concentration in the "biogeochemically coupled" experiment for 1%$\text{yr}^{-1}$ (solid), 2%$\text{yr}^{-1}$ (dashed) and 0.5%$\text{yr}^{-1}$ (dotted) experiments. Panel c) shows the cumulative uptake as a function of temperature in the "radiatively coupled" experiment. Panel d) shows the evolution of the cumulative airborne fraction as a function of the proportional concentration increase for the "fully coupled" experiments.

Figure 7.1 shows the response of the FAIR model (blue) described in section 7.1 under the three experiments described above. The responses of the AR5-IR (red) and BEAM models (green) are also shown. Figure 7.1a shows total carbon uptake by ocean and land, $C_{\mathrm{acc}}$, as a function of time in the 1%yr$^{-1}$ increasing $CO_2$ experiment. As in the full-complexity ESMs shown in Arora *et al.* [2013] and Gregory *et al.* [2009], the coupling between GMST change and the carbon-cycle in the FAIR model acts to suppress carbon uptake, shown by the difference between the thick and thin blue lines, a mechanism that is absent (by construction) in both AR5-IR and biogeochemical version of the BEAM model considered here. The feedback of cumulative carbon uptake on the carbon-cycle in the FAIR model acts to increase airborne fraction in the later stages of the experiment relative to earlier stages (see figure 7.4) and as illustrated by the approximately linear increase in $C_{\mathrm{acc}}$ in the "biogeochemically-coupled" experiment (figure 7.1a).

Figure 7.1b shows $C_{\mathrm{acc}}$ as a function of atmospheric $CO_2$ concentration: again, the FAIR model captures the concave-downward form of this diagnostic show by ESMs, in contrast to the AR5-IR model. Figure 7.1c shows the impact of GMST increase on cumulative uptake, or the difference between the biogeochemically coupled and fully coupled experiments shown in panel (a), as a function of warming. 1%yr$^{-1}$, 0.5%yr$^{-1}$ and 2%yr$^{-1}$ experiments all lie along the same line in panel (c). Panel (d) show cumulative airborne fraction increases after an initial decline, similar to the behaviour of the ESMs. In contrast, the IPCC-AR5 model shows a steady decrease in the cumulative airborne fraction with higher concentrations due to the state-invariant rate at which a pulse of carbon is removed from the atmosphere. The initial cumulative airborne fraction is higher in the AR5-IR model than the FAIR model, as the carbon-cycle response parameters used here are those representative of present-day pulse-response experiments of Joos *et al.* [2013]. An integration of this model under historical emissions therefore produces historical concentrations significantly in excess of those observed in the early 21$^{st}$ century (see figure 7.4a).

The BEAM model (run with parameters as given in Glotter *et al.* [2014], which are tuned to emulate the climate response to $CO_2$ on millennial timescales) displays a very low cumulative uptake of carbon from the atmosphere relative to the other simple models considered here, associated with a cumulative airborne fraction initially around 0.9 which does not decline substantially with time. This model therefore also displays much higher concentrations than observed over the historical period when driven with estimates of historical emissions (not shown), although it must be noted again that the model was tuned to the millennial timescale responses of more complex models and not over the historical period. Although the BEAM model explicitly solves the equations of the ocean carbonate chemistry, it displays a roughly constant cumulative airborne fraction under the exponential concentration increase scenarios considered here, unlike the ESMs which behave more like the FAIR model constructed in this chapter (see figure 3, 4, 5 and 6 in Gregory *et al.* [2009]). In addition, as BEAM focuses solely on oceanic feedback mechanisms it cannot capture saturation of land carbon sinks and the dependence of land carbon sink uptake on warming, an important part of the overall carbon-cycle feedbacks simulated by ESMs [Arora *et al.,* 2013]. Land carbon uptake can contribute an equal (or even greater) fraction of the total system carbon uptake in ESM models as oceanic uptake over time periods on the order of a century. However, it must be noted that physical mechanisms of the land carbon-cycle response to warming remain poorly understood compared with oceanic mechanisms and a wide uncertainty exists about the future land carbon uptake response. Due to this lack of potentially important feedback processes associated with the land carbon-cycle, I do not extend the comparison with the BEAM model in any of the subsequent analyses.

Figure 7.2 shows the response of the simple climate-carbon-cycle models to the emissions pulse experiments of Joos *et al.* [2013]. In these experiments, a set of ESMs are used to derive emissions that are consistent with concentrations rising as observed historically until they exceed 389ppm (achieved in 2012) after which they are subsequently held constant. Figure 7.2a shows the emissions and warming

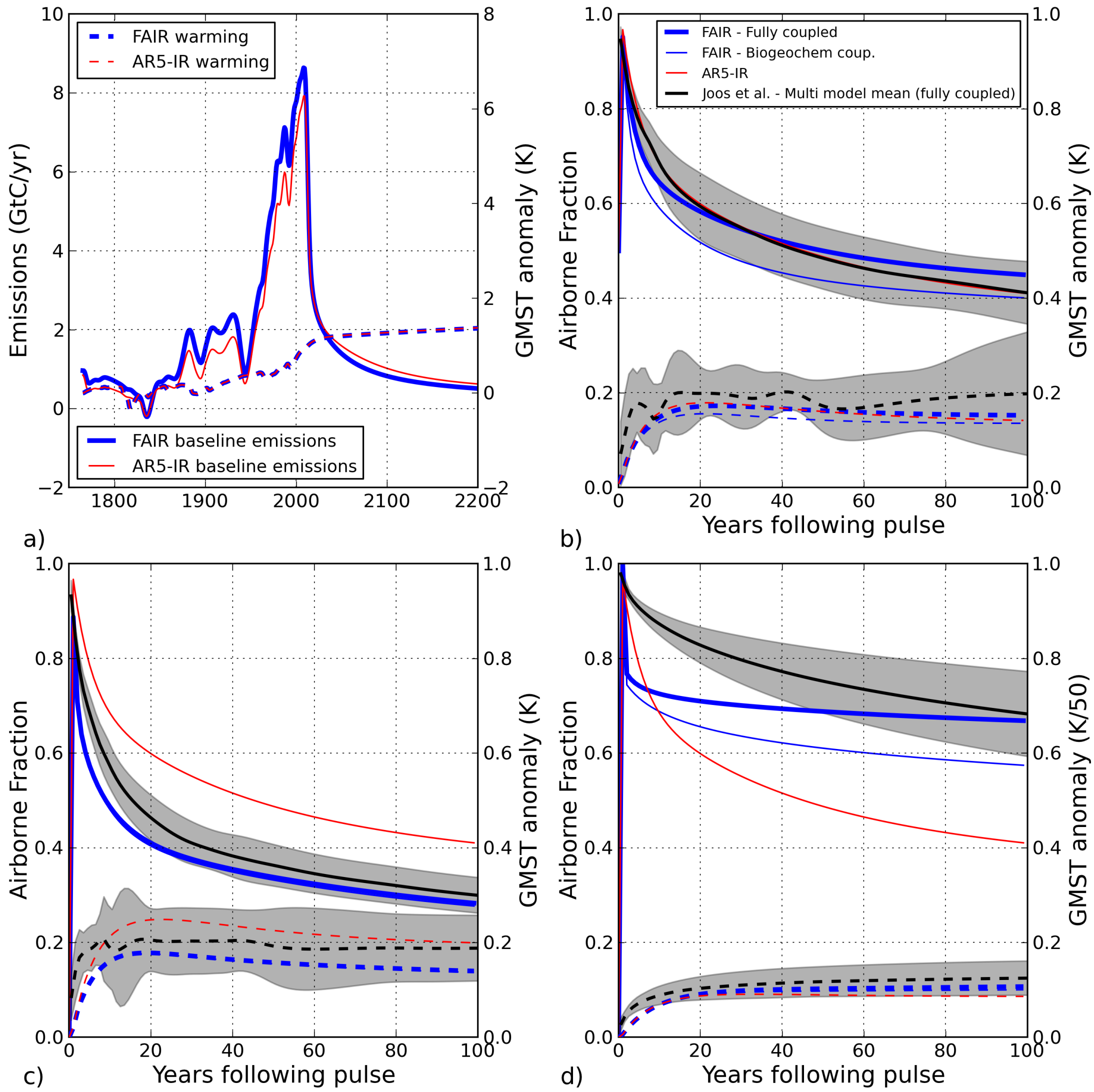


**Figure 7.2** – Response to pulse emission experiments of Joos *et al.* [2013]. Panel a) shows the "baseline" emissions (left-hand axis, solid) and warming (right-hand axis, dashed) when concentrations are stabilised at 389ppm for the FAIR (blue) and AR5-IR (red) models. Panel b) shows the response to a 100GtC imposed on present-day (389ppm) background conditions (PD100 experiment). Panel c) shows the response to a 100GtC pulse in pre-industrial conditions (PI100 experiment). Panel d) shows the response to a 5000GtC pulse in pre-industrial conditions (PI5000 experiment), with the warming normalised by the increase in pulse size between panels c) and d). The black lines in panels b), c) and d) shows the Joos *et al.* [2013] multi-model mean for airborne fraction (solid) and warming (dashed), with the grey shading indicating one standard deviation uncertainty across the ensemble.

compatible with this forcing scenario for the FAIR and AR5-IR models. A declining but sustained low level of diagnosed emissions are required following stabilisation of atmospheric concentrations in order to maintain them at a constant level. In a second experiment, a 100GtC pulse is added to these calculated emissions in the year that concentrations exceed 389ppm and the resulting concentration and temperature anomalies are compared with and without the pulse emission to isolate the coupled response to the pulse emission alone.

Figure 7.2b shows the concentration and warming response to a 100GtC pulse emission emitted against a 389ppm background. After 100 years the concentration anomaly of the pulse emission in the fully coupled FAIR model has decayed to 0.46 of its initial value, slightly greater than the multi-model average of the ESM responses of 0.41, but, the $\mathrm{iIRF}_{100}$ of 53 years is consistent with the ESM multi-model mean of 52.4 years [Joos *et al.,* 2013]. The "biogeochemically-coupled" version of the FAIR model, in which temperature-induced feedbacks are suppressed, reduces the integrated 100-year airborne fraction by 11%. The "fully-coupled" FAIR model shows temperature initially adjusting rapidly followed by near-constant temperature over the remainder of the century, consistent with the ESM behaviour (oscillations are seen in the temperature response of plumes of the ESM and EMIC ensemble in figure 7.2 due to internal variability that isn't simulated by the simple climate-carbon-cycle models).

Figure 7.2c and 7.2d also show the response to a 100GtC and a 5000GtC pulse respectively, applied in pre-industrial conditions. Similarly to the response shown by ESMs discussed in Joos *et al.* [2013], the 100GtC pre-industrial pulse decays faster than the present-day case, due to reduced saturation of the land and ocean carbon sinks. With these parameters, the $\mathrm{iIRF}_{100}$ is approximately 30% lower in the pre-industrial case compared to the present day, consistent with corresponding ratio in the most detailed ESMs, with its value of 36 years within the 34-47 years range of the ESMs. The magnitude of temperature response is similar in both cases due to the increased radiative efficiency of a pulse of $CO_2$ at lower background concentrations counteracting the faster decay of carbon out of the atmosphere. The 89% increase of $iIRF_{100}$ in the 5000GtC pre-industrial pulse relative to the 100GtC pulse against a pre-industrial background shows, whilst smaller than the approximate doubling observed in the ESMs, that the FAIR model can capture the dependence of the pulse-response on pulse size as well as background conditions, whilst the AR5-IR model displays, by construction, identical pulse response independent of pulse size or background conditions.

A difference between the FAIR model and the ESMs is that restricting temperature-induced feedbacks on the carbon-cycle does not result in a substantial reduction in the $iIRF_{100}$ for the pre-industrial 100GtC pulse experiment (the "fully-coupled" and "biogeochemically-coupled" experiments lie on top of each other in figure 7.2c), whereas a 13% reduction in $iIRF_{100}$ is observed for the ESMs [Joos *et al.,* 2013]. It is only for the 5000GtC pre-industrial pulse experiment that a reduction is seen in the FAIR simulated $iIRF_{100}$ due to suppression of the temperature-induced feedbacks on the carbon-cycle.

Herrington & Zickfeld [2014] conducted several idealised fixed cumulative emissions experiments with the UVic Earth System Model of intermediate complexity [Weaver *et al.,* 2001]. I here emulate the PULSE experiments of Herrington & Zickfeld [2014] by integrating the FAIR model with historical fossil fuel and land-use $CO_2$ emissions (as derived from historical concentrations using the MAGICC model [Meinshausen *et al.,* 2011]) together with estimates of the historical radiative forcing from non-$CO_2$ factors. Pulse emissions of various sizes were then applied over a two-year period from 2008 in order to restrict total all time cumulative emissions to specified totals (see Herrington & Zickfeld [2014] for details). Non-$CO_2$ forcings are held constant at 2008 levels after following RCP8.5 [Riahi *et al.,* 2011] trajectories for 2005-2008.

Figure 7.3 shows the response of the FAIR model, as well as the AR5-IR model, to the experiments described above. For all pulse sizes (denoted with different

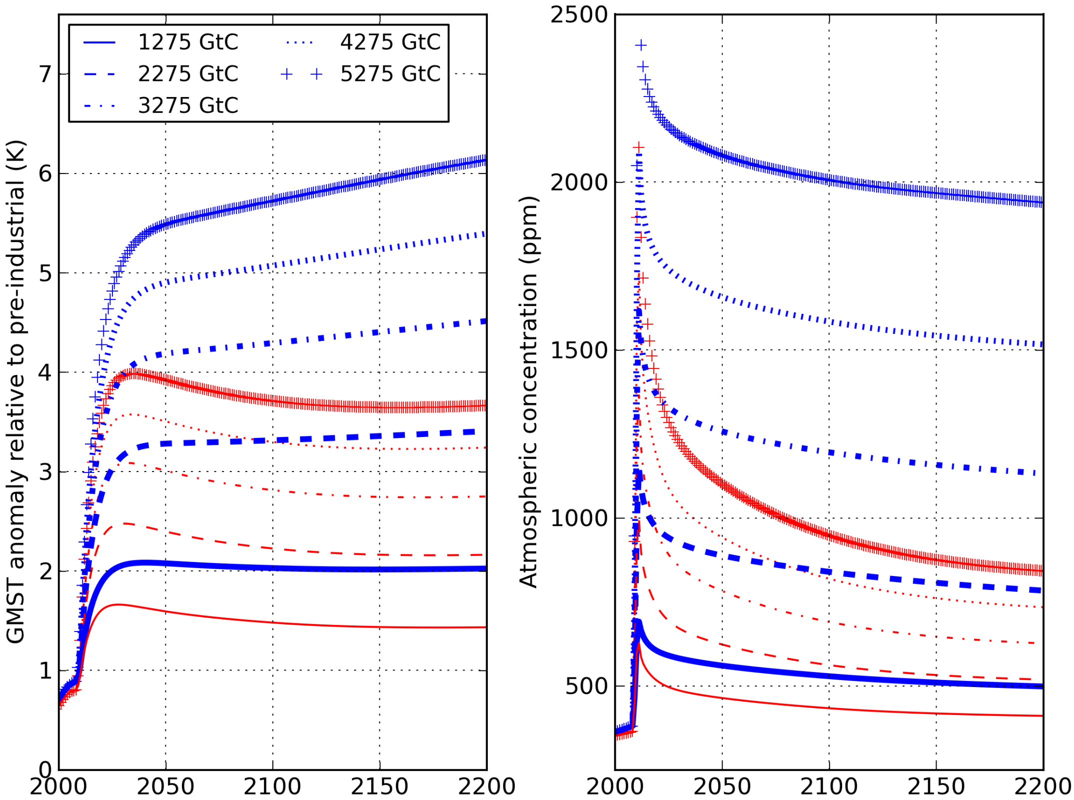


**Figure 7.3** – The left hand panel shows the GMST response to the pulse experiments of Herrington & Zickfeld [2014]. Pulse emissions are applied over a 2-year period from 2008, with differing total cumulative carbon emissions denoted by different line styles. Responses are shown for the FAIR (blue) and AR5-IR (red) models. The right hand panel shows the corresponding concentration response.

linestyles) contrasting the fully coupled FAIR (thick blue) and the AR5-IR (red) models shows that including carbon-cycle feedbacks is essential to prevent a substantial decay in the temperature anomaly over the first 100 years following the pulse emission. Ricke & Caldeira [2014] used a version of the AR5-IR model to find that the maximum warming from a pulse emissions of $CO_2$ occurs approximately a decade after emission, but as shown here and as highlighted by Zickfeld & Herrington [2015], not accounting for feedbacks on the carbon-cycle fails to capture the dependence of the shape of the plateau of $CO_2$-induced warming over the century following emission cessation and the timing of peak warming on the size of the emissions pulse. At higher pulse sizes, the temperature response in the FAIR model fails to plateau as quickly as at lower pulses, where the balance between carbon-cycle cooling and long-timescale thermal warming takes centuries to reach balance (similarly as shown for UVic in figure 3 of Herrington & Zickfeld [2014]).

As a validation of the FAIR model over the historical period, figure 7.4 shows the comparisons to historical data from Le Quéré *et al.* [2015]. In figure 7.4 simulations are commenced from 1850, which is assumed to be a quasi-equilibrium state for the carbon-cycle, in order to facilitate comparisons with the observed data (which is available from 1850 onwards). Panel (a) shows the concentration response to estimated global historical $CO_2$ emissions. The FAIR model replicates concentrations over the past several decades well, as opposed to the AR5-IR model, which has a bias of over 30ppm in 2011. Similarly, emissions derived from the time series of historical atmospheric $CO_2$ concentrations (figure 7.4b), are lower in AR5-IR model, due to the slower decay of $CO_2$ from the atmosphere over the historical period. Whilst the time mean-value of the airborne fraction (defined as the fraction of emissions remaining in the air after one year) is captured well by the FAIR model, fluctuations are of much greater magnitude in the observed record, indicating short timescale processes and variability in the carbon-cycle that is not captured by these simple models. The AR5-IR (which is tuned to the present-day response) displays a too large airborne fraction over the entire historical period and is less consistent

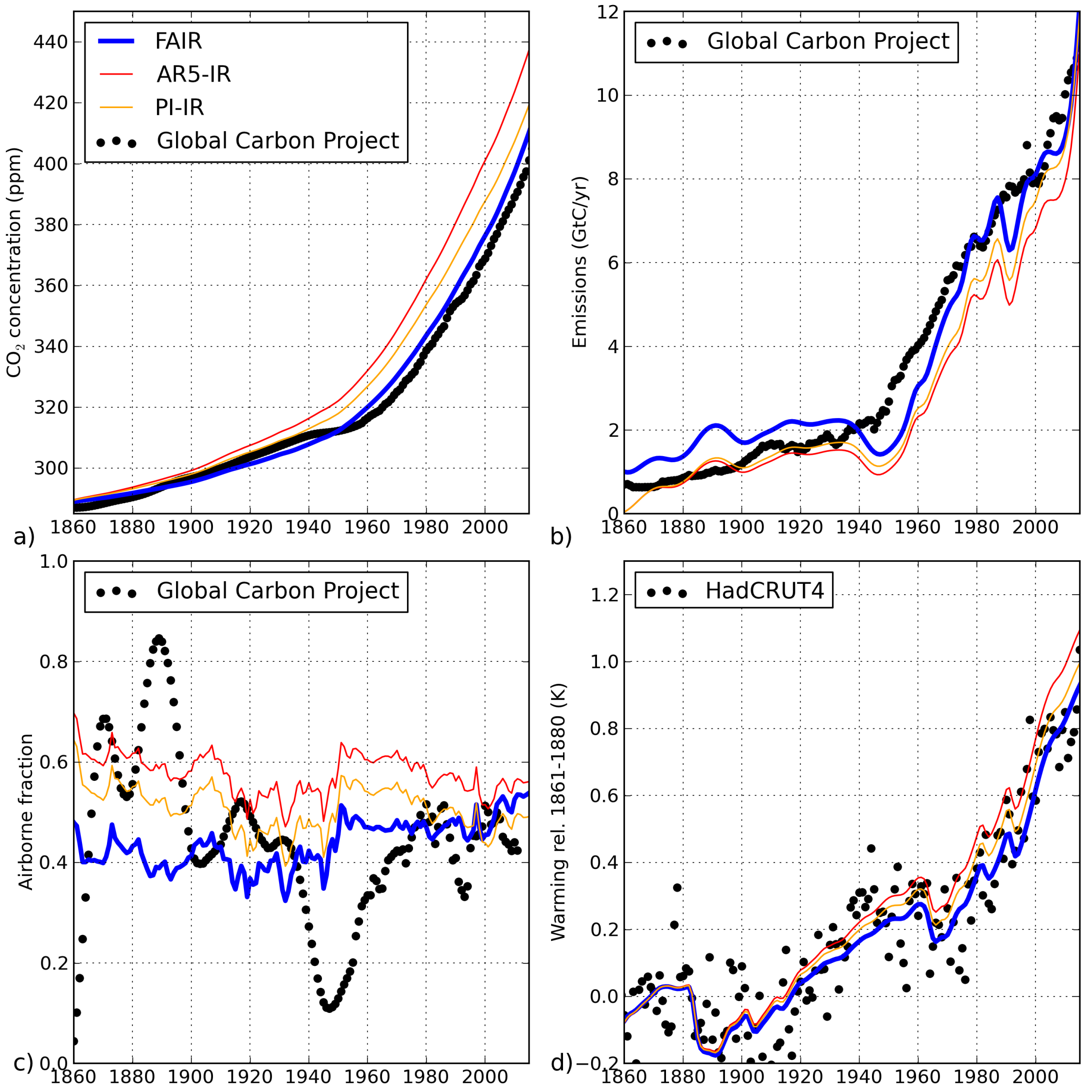


**Figure 7.4** – Panel a) shows the $CO_2$ concentration response when integrated under historical emissions (and non-$CO_2$ radiative forcing) for the FAIR (blue), AR5-IR (red), and PI-IR (orange) models. Panel b) shows the derived $CO_2$ emissions consistent with historical concentrations. Panel c) shows the evolution of annual airborne fraction (smoothed with a 7-year running mean for the observations) in the models when driven by historical emissions. Panel d) shows the warming anomaly in the models when driven by historical emissions. Historical observations are shown as black dots in all panels. Panels a), b) and c) all show data from Le Quéré *et al.* [2015] and panel d) shows the HadCRUT4 [Morice *et al.*, 2012] dataset.

with the observations than the FAIR model (figure 7.4d). Retuning the AR5-IR model to pre-industrial conditions (PI-IR - orange line in figure 7.4 and figure 7.5, with coefficients derived from table S2 of Joos *et al.* [2013]) would improve its fit to past emissions, but would then give a worse fit to the behaviour of more complex models under impulse-response experiments and future scenarios (see figure 7.5). Historical temperature changes are simulated similarly under all versions of the carbon cycle model, but temperatures are greatest in the AR5-IR case where the above observed $CO_2$ concentrations result in warming that is greater than the HadCRUT4 observations over the recent decade.

I also check the FAIR model's response to a pair of benchmark scenarios against that of the widely used MAGICC model [Meinshausen *et al.,* 2011], a simplified box-model of the climate system response to emissions of various greenhouse gases that has commonly been used to assess climate mitigation scenarios [Clarke, 2014]. Panels a) and b) of figure 7.5 show the $CO_2$ concentration from the FAIR model under RCP8.5 and RCP2.6 emissions respectively when forced with emissions diagnosed from the RCP concentration profiles using MAGICC, contrasted with the concentration timeseries from MAGICC (green), which is by definition equal to the corresponding scenario-defined concentration profiles. Non-$CO_2$ forcing is the same in all cases.

The FAIR model compares well to MAGICC, particularly for the ambitious mitigation scenario (which is less well reproduced by other simplified climate-carbon-cycle models such as BEAM). There is some divergence after 2100 in the high emission scenario, but the behaviour of MAGICC (or indeed any other model) under these more extreme forcing scenarios has not been verified. Whilst comparing the performance of one simple model to another is not as rigorous a test of model performance as comparing directly to the behaviour of ESMs, it is encouraging that the FAIR model shows a close correspondence with a well-known and well-used simple model that has been used extensively to emulate the response of ESMs [Rogelj *et al.,* 2012].

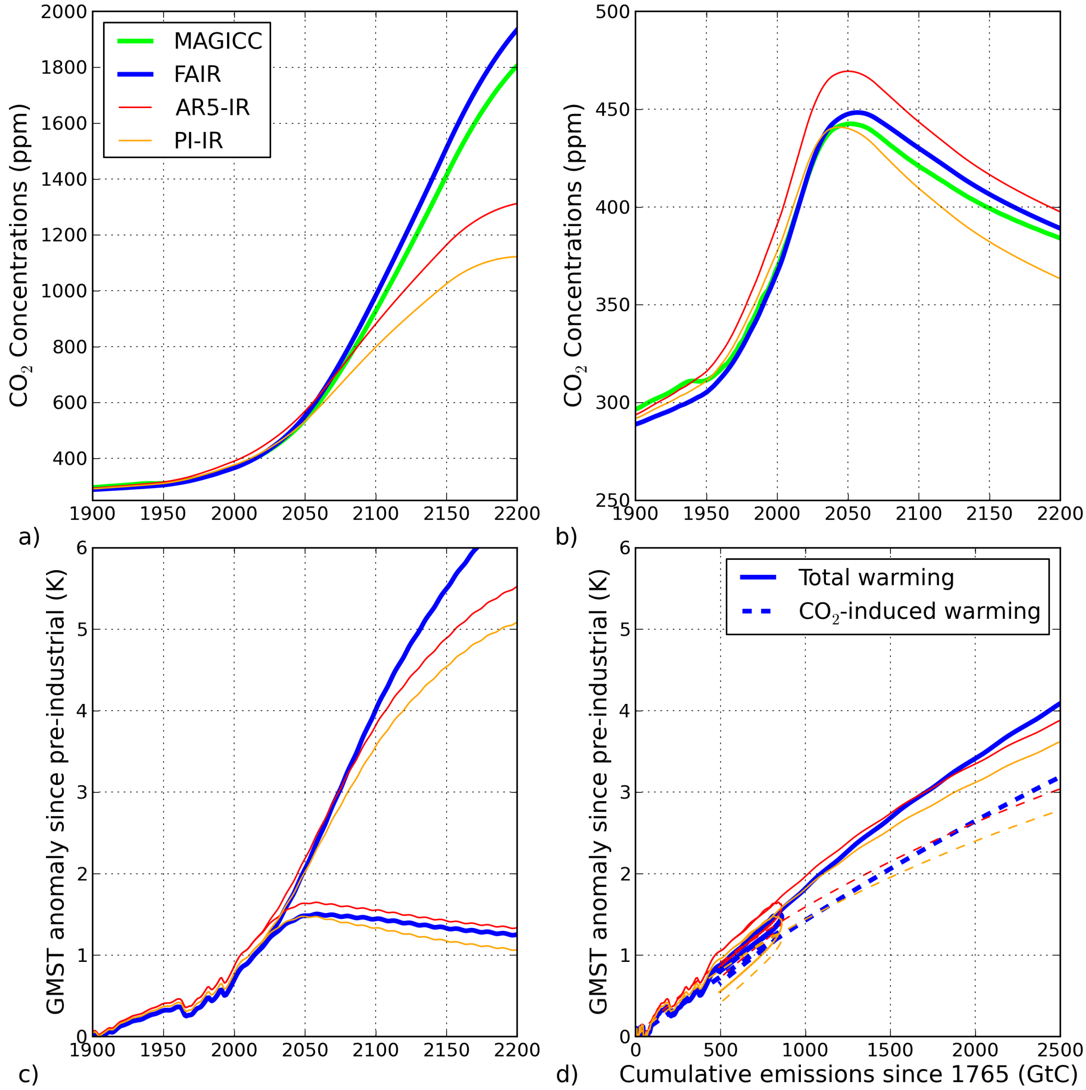


**Figure 7.5** – Panels a) and b) shows the $CO_2$ concentrations under RCP8.5 and RCP2.6 respectively for the FAIR (blue), AR5-IR (red), PI-IR (orange) and MAGICC (green) models. Panel c) shows the temperature response under both RCP2.6 and RCP8.5. Panel d) shows the evolution of total warming (full) and $CO_2$-induced warming (dashed) as a function of cumulative carbon emissions.

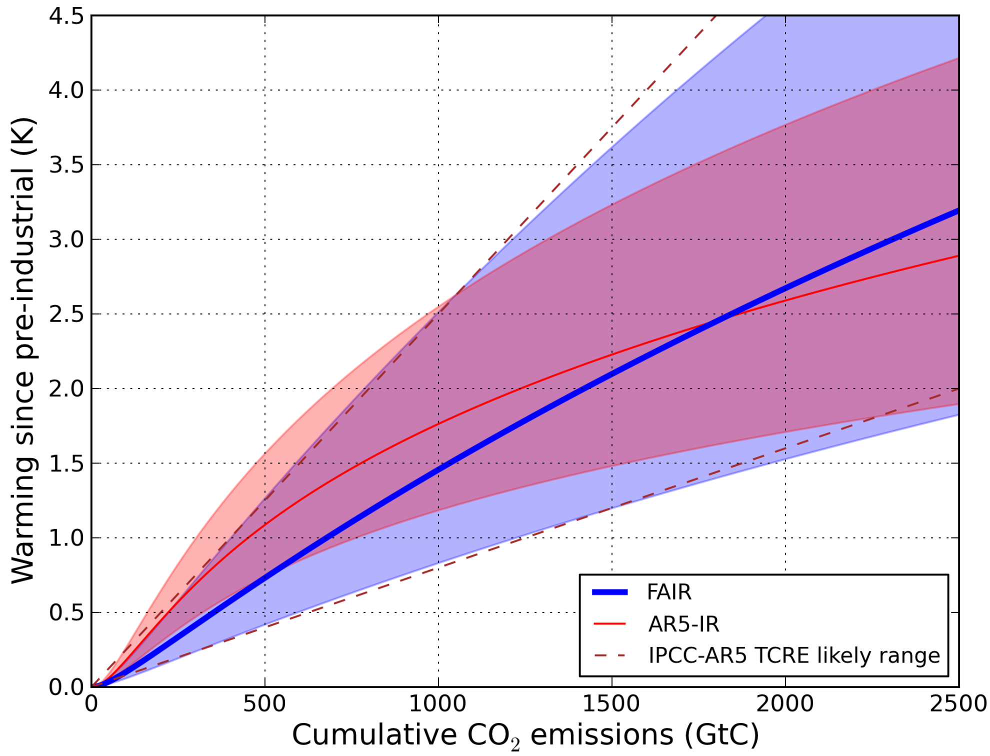


**Figure 7.6** – TCRE uncertainty in the FAIR model. Straight lines indicate a constant TCRE. Dashed brown lines show the IPCC-AR5 *likely* 0.8-2.5K/TtC assessed range for TCRE. The blue plume shows the response to 1%$yr^{-1}$ increase in $CO_2$ concentrations for the IPCC-AR5 *likely* TCR and ECS ranges in the FAIR model, with an additional +/-10% perturbation to the $r_0$, $r_T$ and $r_C$ parameters for the high/low end the *likely* ranges respectively. The red plume shows the AR5-IR model response.

Figure 7.5c shows the temperature responses to the RCP8.5 and RCP2.6 scenarios, and the relationship between temperature and cumulative carbon emissions for the same scenarios is shown in figure 7.5d. Crucially, the FAIR model captures the straight-line relationship between cumulative carbon emissions and human-induced warming that was highlighted in the IPCC-AR5, and is becoming an integral part of climate change policy analysis [Knutti & Rogelj, 2015; Matthews, 2015; Millar *et al.,* 2016; Rogelj *et al.,* 2016]. When integrated under a 1%yr$^{-1}$ concentration increase scenario, the FAIR model, with parameter settings given in section 2, has a TCRE =1.5K/TtC (see figure 7.6). A common +/- 10% perturbation to the parameters $r_T$ , $r_C$ and $r_0$ (combined with perturbations to $c_1$ and $c_2$ consistent with the IPCC-AR5 *likely* ranges for TCR: 1.0-2.5K and ECS: 1.5-4.5K) allow the IPCC-AR5 *likely* TCRE range of 0.8-2.5K/TtC to be spanned, as shown in figure 7.6. In contrast, the AR5-IR model, with a constant airborne fraction, shows a clear concave-downward shape in a plot of realised warming against cumulative carbon emissions, because the decline of the cumulative airborne fraction is unable to compensate (as it does in more complex models) for the logarithmic relationship between $CO_2$ concentration and radiative forcing [Caldeira & Kasting, 1992; Millar *et al.,* 2016]. The FAIR model also displays some curvature at high cumulative emissions, consistent with the behaviour for ESMs described by Leduc *et al.* [2015].

## 7.2 An adaptive policy for long-lived climate pollutants

Thus far in this chapter, I have developed and validated a simple climate-carbon-cycle model, FAIR, that captures key physical-dependencies of the response of ESMs under idealised emission scenarios, whilst maintaining a simple and transparent structure for sampling physical climate response. I here use this model to develop a new 'adaptive' mitigation framework that can be robust to the broad uncertainty in physical climate response found in the rest of this thesis.

The stated goal of international climate policy is to limit GMST warming to "well below 2°C" above pre-industrial levels, and to pursue efforts to meet a 1.5°C goal [UNFCCC, 2015]. Whether or not these temperature limit are economically optimal (which is beyond the scope of this thesis), the scientific consensus concludes that risks of dangerous anthropogenic interference in the climate system increase systematically at higher levels of warming [IPCC, 2014a]. As described in section 7.1.1, even after emissions cease temperatures will remain largely flat (except for natural variations) for decades to centuries [Solomon *et al.,* 2009], locking in permanent climate change on human relevant timescales. It has been argued [Pierrehumbert, 2014] that this permanence justifies a continued focus on limiting overall $CO_2$-induced warming, which is primarily driven by cumulative $CO_2$ emissions [Allen *et al.,* 2009b; Meinshausen *et al.,* 2009], rather than a focus on shorter-term policy targets at specific dates [Victor & Kennel, 2014], which may act to marginally reduce warming in the short-term.

Cumulative emission budgets are subject to both physical and socio-economic uncertainties. Cumulative $CO_2$ emissions need to be limited to less than 790 GtC (almost 600GtC of which have already been emitted) in order to give a better than two in three chance of preventing peak GMST rise of more than 2K under standard assumptions for the climate response and future emissions of non-$CO_2$ climate pollutants [IPCC, 2013]. If, however, net non-$CO_2$ warming can be reduced by 0.5K by the time temperatures peak [Shindell *et al.,* 2012], the corresponding cumulative $CO_2$ emission budget rises to 1000 GtC, or 1210 GtC for an even chance of exceeding 2K: a three-fold uncertainty in remaining allowable $CO_2$ emissions. Constructing a global climate policy to limit warming to less than 2K under these uncertainties without either exceeding the actual carbon budget or undergoing potentially costly overly-rapid mitigation is clearly a challenge.

Future warming is, to first order, a function of future emissions only [Matthews &

Caldeira, 2008]. Hence, in order to prevent a global temperature rise of more than 2K, global net emissions of $CO_2$ (and other long-lived greenhouse gases) must fall to zero by the time warming reaches 2K. Uncertainty in the TCRE, such as the kind discussed in the previous chapters of this thesis (via uncertainty in TCR and ECS), makes designing a fixed future emissions pathway that is both economically feasible and will achieve a 2K or less peak warming a difficult task, due to the range of cumulative emissions that could be consistent with the temperature goal.

### 7.2.1 Economic and policy characteristics of response independent 2K mitigation pathways

As examples of 'climate response independent' (where global emissions evolve independent of climate change) emissions scenarios, figure 7.7a (non-blue curves in this panel and all other panels) shows global emissions scenarios from Working Group 3 (WG3) and Synthesis Reports of the IPCC-AR5 [Clarke, 2014; IPCC, 2014b; IPCC, 2014c]. These scenarios are 'policy-optimal' (in the sense that all countries of the world begin mitigation immediately, there is a single global carbon price, and all key technologies are available [IPCC, 2014b]) energy transition trajectories that have cumulative $CO_2$ emissions roughly consistent with 2100 temperatures being *likely* less than 2K above pre-industrial levels. This subset of scenarios was used to assess the cost of the climate mitigation required to keep 2100 temperatures beneath 2K and are an indicative set of economically-efficient climate response independent emission trajectories that target achieving a specific atmosphere $CO_2$ concentration in 2100.

These scenarios are derived from integrated assessment models (IAMs) as optimal pathways to minimise the global-aggregate economic cost of achieving the 2K carbon budget, under the constraints and assumptions of a particular IAM, including the availability and cost of various mitigation technologies. Even under the idealised policy-optimal assumptions, the overall mitigation costs can be high (as

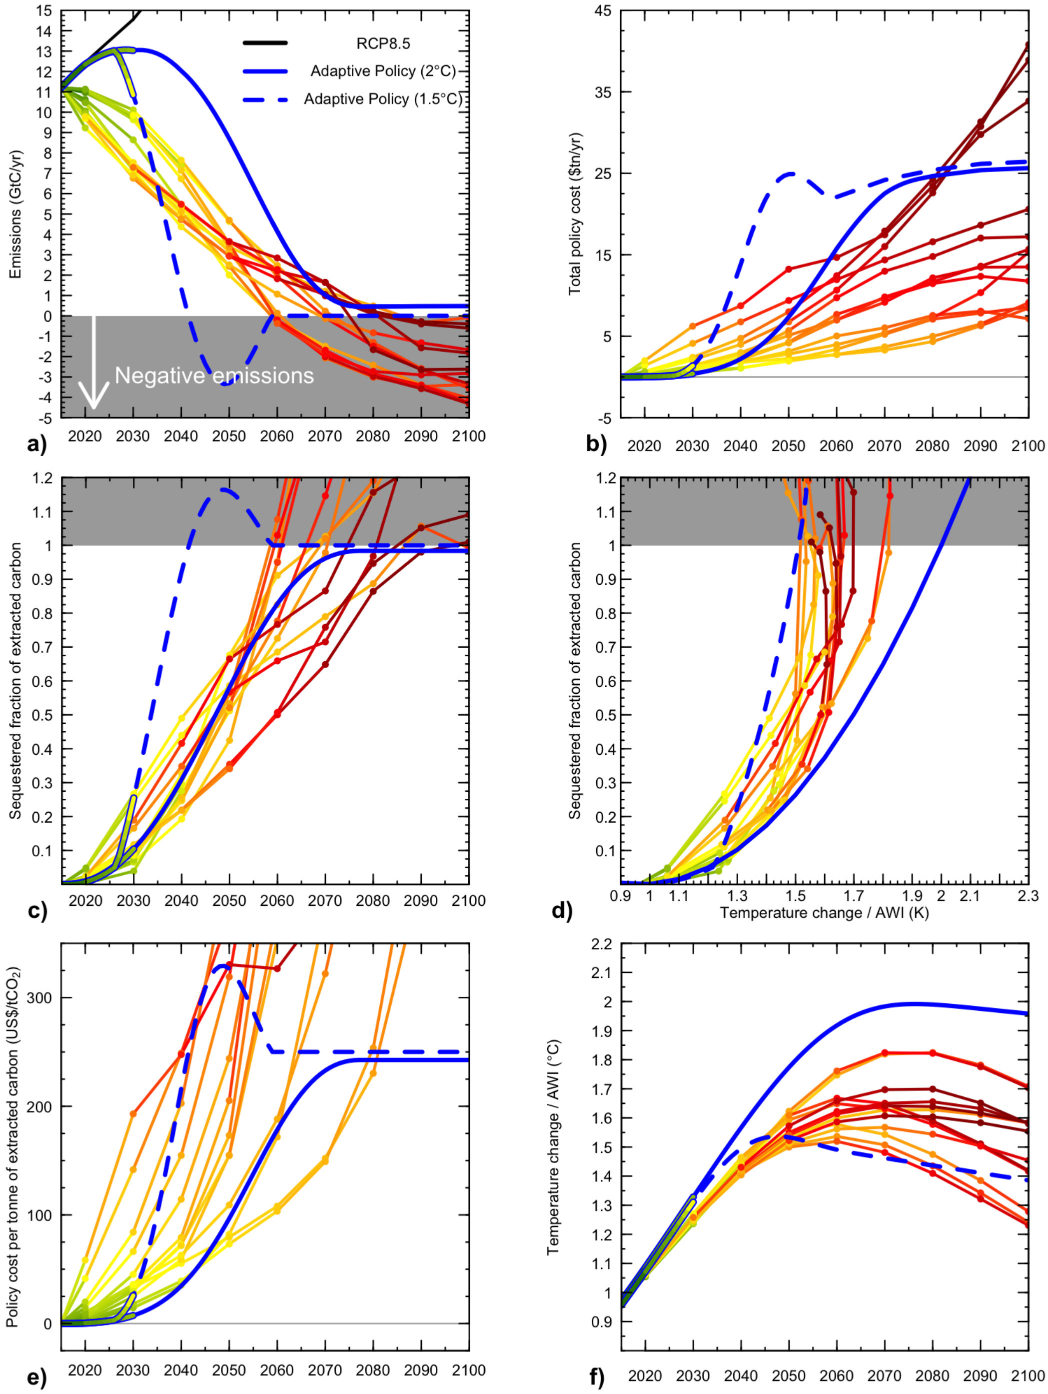


**Figure 7.7** – "Cost-optimal" scenarios that respect the cumulative emissions budget consistent with a 2K maximum warming compared with a hypothetical sequestration-based mitigation scenario. In all panels, cost-optimal scenarios are coloured by the total annual policy cost (consumption loss relative to a baseline scenario) of mitigation policies as computed by the individual IAMs: a) annual $CO_2$ emissions; b) total policy cost in 2005 US$; c) sequestered fraction of extracted carbon, S, as a function of time; d) S as a function of the attributable warming index (AWI); e) policy cost per tonne of extracted carbon, $C_e$, as function of time; and f) 21st century GMST (equivalent to the AWI) as a function of time. Black line in panel a) shows the RCP8.5 baseline emissions. The solid blue curves represent a 2K mandatory sequestration scenario defined in panel d). The dashed blue curve represents a scenario in which the target of the adapative policy is changed from 2K to 1.5K at an AWI of 1.2K. In all panels apart from panel d) the blue curves are coloured by total policy cost out to 2030 based on consumption of carbon continuing to follow RCP8.5.

show by the range of mitigation policy-costs relative to baseline shown in figure 7.7b). To reduce the impact of difference in reported emissions between 2010 (the policy start dates in many scenarios) and today, I 'harmonise' the emissions scenarios to match RCP8.5 emissions until 2015. I use scaling factors to harmonise the emissions timeseries that linearly decline to zero in 2050 in order to preserve as much information from the original scenario as possible, particularly in the second half of the century when peak warming is likely to be reached [Rogelj *et al.,* 2011]. Mathematically, the harmonised emissions, $E_{harm}(t)$, can be described as,

$$\begin{aligned} 2015 \leq t < 2050 : E_{harm}(t) &= E(t)\left(\frac{E(2015)}{E_R(2015)} + \left[1 - \frac{E(2015)}{E_R(2015)}\right]\left[\frac{t-2015}{2050-2015}\right]\right) \\ t \geq 2050 : E_{harm}(t) &= E(t), \end{aligned} \tag{7.2.1}$$

where $E(t)$ is the spline-smoothed (in order to produce an annual timeseries from the scenario data reported at 10 year time intervals) unharmonised scenario emissions, and $E_R(t)$ the reference RCP8.5 emissions.

The need for net negative emissions technologies, such as biomass energy combined with CCS (BECCS), by the end of the century is clear in many scenarios (figure 7.7a). Figure 7.7a also contrasts the scale of the emission cuts that are needed to remain within the *likely* <2K carbon budget and a 'no climate policy alternative (RCP8.5, black line).

Despite the diversity of the shapes of the emissions scenarios, the scenarios look more similar when expressed in terms of variables that are physically important for the global climate response. Figure 7.7c shows the evolution of the fraction of extracted carbon that is sequestered, S, over time in these scenarios. S is defined as the ratio of net carbon sequestered in a given year, accounting for leakage, to total carbon consumed (both emitted and sequestered), in the same year. Because of the cumulative impact of $CO_2$ emissions, and the expected continued usage of some amount of fossil carbon for certain economic activities at the end of the century in all mitigation scenarios, S must eventually reach unity to ultimately limit warming. If committing to solar radiative management geoengineering measures to permanently restrict warming are to be avoided, S must eventually reach unity, irrespective of whether stabilisation is at 2K, 3K or 4K. This explains why all scenarios (which all succeed in limiting GMST warming to some finite value) display approximately the same shaped pathway in figure 7.7c.

Figure 7.7d shows the evolution of S in relation to the GMST change calculated by integration of the emissions scenarios with the simple climate-carbon-cycle model (FAIR) described in section 7.1 under a mid-range climate response (TCR=1.6K, ECS=2.75K and a conservative estimate of the carbon-cycle parameters: $r_0$=35.5yr, $r_T$=6.5yr/K and $r_C$=0.035yr/GtC). These values for the climate response parameters are chosen in order to match the median physical climate response of the MAGICC model used in Meinshausen *et al.* [2009] and Rogelj *et al.* [2012] and give a TCRE of 1.59K/TtC. In all scenarios the non-$CO_2$ greenhouse gas radiative forcing is derived from the scenario data where reported or the mean of all reported scenarios when not. The natural radiative forcing is smoothed with a 10-year Gaussian filter after 2015.

All the scenarios in figure 7.7c represent emissions scenarios with a *likely* chance of limiting peak warming to beneath 2K. In these scenarios, S reaches unity when temperatures reach 1.4-1.7 K under the best-estimate climate response (consistent with a better than two in three chance of limiting peak warming to <2K, accounting for uncertainty in the climate response), followed in most cases by a period where S>1 (net negative emissions). Net negative emissions allow temperatures to be reduced from their peak value, but peak warming itself is a function of the cumulative emissions to the time they reach net zero.

### 7.2.2 Constructing a response-dependent $CO_2$ mitigation policy to meet a 2K goal

Section 7.2.1 demonstrated that the evolution of the fraction of extracted carbon sequestered, S, is a physically important policy variable to the value of peak warming. S is a policy variable that indicates the ability to ultimately limit the cumulative stock of $CO_2$ in the atmosphere, which will be required to ultimately limit warming to any temperature. S would therefore seem a sensible and physically-meaningful variable through which to describe an emissions mitigation policy.

Indexing increases in S with cumulative $CO_2$ emissions, such that S=1 after a certain carbon budget has been expended, is still fragile to uncertainty in the physical climate response. As described above, uncertain assumptions about the contribution of non-$CO_2$ warming and the currently irreducible uncertainty in the physical response of the climate to $CO_2$ emissions (see the previous chapters of this thesis), create a broad range of possible cumulative emissions budgets for a 2K peak warming threshold. Indexing increases in S against carbon budgets therefore creates a risk of permanently over-shooting a temperature goal for an overestimate of the remaining carbon budget (or alternatively potentially costly over-mitigation for an underestimate of the remaining budget) and is therefore still highly susceptible to inherent and significant uncertainty in the physical climate response. This section describes an alternative policy formulation in which the schedule for increasing S is indexed to the observed evolution of the physical climate state, aiming to decouple risks of overshooting peak warming targets from uncertainty in the physical response of the climate system.

Otto *et al.* [2015] described an index of observed anthropogenic warming (the AWI-defined as the component of observed warming at any time that is attributable to humanity's impact on the climate). This index is calculated in Otto *et al.* [2015] using the simple model used in chapter 4, but can in theory be calculated using more complex GCMs via the established detection and attribution framework [Bindoff, 2013]. For the set of IPCC-WG3 scenarios shown in figure 7.7d, the GMST anomaly (x-axis) is a good indicator of the AWI across the rest of the century as natural forcing has here assumed to be zero in the future (except for the solar cycle) and natural variability is not captured within the model.

The evolution of the sequestered fraction of extracted carbon, S, as a function of the AWI (figure 7.7d) immediately suggests a convenient relationship for a carbon policy based on S: almost all the "likely 2K" scenarios have a similar shape, curving upwards monotonically with rising cumulative emissions to reach S=1 when the AWI=1.5-1.8K, followed in most cases by a period where S>1 (net negative emissions). IPCC-WG3 scenarios meeting higher budgets and temperature goals all have the same shape (not shown). Based on this, the blue curve in figure 7.7d shows a proposed policy based on increasing S with AWI increases. The simplest form of this policy is a quadratically rising S (from S=0 at AWI=0.93K, indicative of attributable warming to-date [Otto *et al.,* 2015]) as the AWI increases towards the global temperature goal, $\mathrm{AWI}_{\mathrm{max}}$ (here assumed to be 2K) and reaching unity when the AWI reaches $\mathrm{AWI}_{\mathrm{max}}$. A quadratic, as opposed to linear, function is chosen to allow the initial rate of increase in S to be close to zero (as it is today). The function form of the adaptive policy can hence be written down as,

$$\mathrm{S}(\mathrm{AWI}(\mathrm{t})) = \left(\frac{\mathrm{AWI}(\mathrm{t}) - \mathrm{AWI}_{2015}}{\mathrm{AWI}_{\mathrm{max}} - \mathrm{AWI}_{2015}}\right)^2. \qquad (7.2.2)$$

As discussed above, an alternative policy would be to link the increase in S to the emitted cumulative carbon emissions. However, such a policy would require an apriori choice about the size of the future cumulative carbon budget allowable, and success in achieving a temperature goal would be dependent on the climate response, unlike policies explicitly linked to the AWI.

Corresponding emissions for this proposed policy are shown by the blue curve in figure 7.7a. This scenario assumes non-$\mathrm{CO}_2$ radiative forcing follows the mean of the 13 IPCC scenarios that report the sufficient breakdown in radiative forcing. In the AWI-indexed policy, using the best-estimate of the climate response, temperatures peak just before emissions reach zero (figure 7.7d), after which S is assumed constant. Such allowable 'emission floors' are very uncertain [Bowerman *et al.,* 2011], and for the purposes of todays decisions, are effectively indistinguishable from zero. In this policy proposal, changes in the sequestered fraction are explicitly linked to the human attributable warming through the AWI and hence respond to the identifiable long-lasting warming signal and not to short-lived natural fluctuations. Whilst the IAM scenarios displayed are consistent with a better than two in three chance of limiting peak warming to 2K, they give a peak warming significantly less than 2K under a best-estimate response. The adaptive policy scenario, with an explicit target of 2K, can allow significantly higher emissions in the near term because it is designed to tighten automatically as the actual climate response emerges.

The dashed blue curve in the panels of figure 7.7 shows a scenario in which the temperature goal of the adaptive policy is changed to 1.5K when temperatures reach 1.2K (reached around 2025-2030). This corresponds to a firm commitment to achieve the higher-ambition goal of the Paris Agreement, which may occur if emergent climate impacts are greater than currently expected. In such a scenario, under the adaptive climate policy emissions would be required to fall rapidly to zero in about 15 years, however, a small overshoot of the 1.5K target would still be unavoidable. Both the 2K and 1.5K adaptive climate policy scenarios shown here display large peak emissions reduction rates which may not be economically sustainable, and in part reflect the very high carbon baseline (RCP8.5) used in these scenarios. Additional near-term mitigation actions are an essential complement to an adaptive sequestration-based policy in order to reduce the risk of requiring unsustainable rates of decarbonisation later in the century.

Thus far I have assumed a fixed climate response temperature in calculating the GMST changes for the emissions scenarios shown in figure 7.7. The advantage of an adaptive sequestration policy indexed to AWI is most evident when the climate

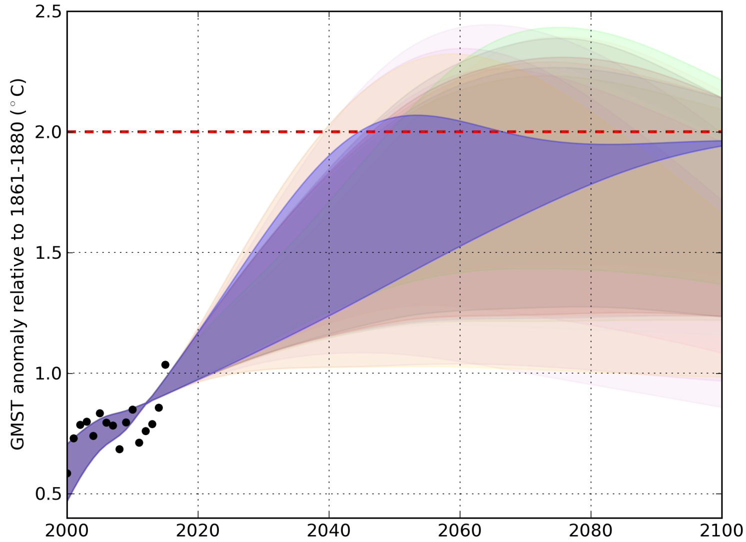


**Figure 7.8** – GMST anomaly (relative to 1861-80) for the IAM scenarios from figure 7.7a (non-blue plumes) under a sampling of physical climate response uncertainty. The plumes denote the minimum and maximum of the temperature change in each emissions scenario at any point in time under the climate response uncertainty sampling. The blue plume denotes the spread of the adaptive sequestration-based policy under the same climate response uncertainty sampling, with a 2K temperature target (red dashed line) The black dots show the HadCRUT4 observations.

response is uncertain. Throughout the rest of this thesis endemic and currently irreducible uncertainty in the climate response has been documented. Developing climate emissions targets independently of the observed climate response (such as any of the non-blue curves in figure 7.7) will therefore leave the peak temperature anomaly (and the scale of climate impacts) uncertain. Figure 7.8 shows the same scenarios as in figure 7.7 integrated under a sampling of different conceivable physical climate responses. For all scenarios, I independently sample TCR, RWF within the ranges 1.0-2.5K and 0.55-0.66 respectively, corresponding to a TCRE range of 1.0-2.6K/TtC. Anthropogenic aerosol radiative forcing is scaled as described in section 4.2 in order to achieve approximate consistency with the present-day state of the climate system.

The IPCC-WG3 IAM scenarios (green to pink plumes) show a large range of warming by 2100. In contrast the range of responses to this adaptive sequestration-based policy, shown by the blue plume, displays much less sensitivity to the climate response because of the inherent negative feedback built into the policy. The sequestered fraction in the adaptive policy scenarios is no longer constrained to be constant after a peak in temperature, unlike in figure 7.7d. Even with a high climate response (TCRE=2.6K/TtC), the maximum temperature for this policy scenario barely rises above 2K. Similarly, if the observed climate response turns out to be on the low side, then the economic gains of emitting the extra carbon can be realised for longer to take advantage of the lower-than-expected climate damages.

Crucially, the adaptive policy uses only observed warming in the calculation of the AWI index: at no point does the policy require any foresight either of the climate response to emissions or the policy interventions required to reduce emissions to zero. Transparency is a key advantage of this approach: such an adaptive policy would demonstrate a clear and up-to-date link between mitigation actions and what is required to achieve the ultimate goal of stabilising temperatures. Indexing policies against attributable human-induced warming reduces the volatility of the index relative to indexing against the raw observed global temperatures, which include natural variability. The AWI can be easily up-dated with new observations of the climate system as they are made due to the simplicity of the methodology and hence avoid delays in the understanding the true current anthropogenic warming of the climate system.

Some uncertainties may arise in the implementation of an AWI indexed approach due to the ambiguous definition of global temperature in the the Paris Agreement. Richardson *et al.* [2016] showed that a true global mean present-day warming in excess of 1°C could be reconciled with the HadCRUT4 observations, in part due to the observational mask and the blending of land and ocean temperatures to derive the observational record. However, the relevance of unmeasurable temperature changes for assessing the present-day state of the climate in a policy-relevant context is unclear. Additionally, we have here assumed that there is no delay between diagnosing the state of the climate system and implementing mitigation actions in response. In reality there are likely to be socio-economic lags to how fast the sequestered fraction of carbon can respond to changes in the climate state. In particular, the Paris Agreement pledged five yearly reviews of mitigation commitments, suggesting that updates of mitigation commitments based on the evolution AWI is only likely to be possible with this regularity. However, a regular update of mitigation commitments in light of realised climate change in order to meet the long-term goals of the Paris Agreement would be very much in the spirit of the proposed regular 'pledge and review' framework [UNFCCC, 2015] for future increases in climate mitigation ambition.

## 7.3 Economic implications for adaptive climate policies

A similar policy based on an increasing fraction of extracted carbon sequestered was proposed by Allen *et al.* [2009a] (without response-independent indexing to the AWI), but no assessment was made of the cost relative to more conventional mitigation options. To explore the cost considerations of the policy proposal in section 7.2.2, I here assume that, for low values of the sequestered fraction, the cost of sequestering a tonne of $CO_2$, $C_s$, is 50\$ [Dahowski *et al.,* 2012]. This is based on the assumption that sequestration through power sector carbon capture and storage (CCS) will initially be deployed as part of economically advantageous activities, such as enhanced oil recovery [Metz *et al.,* 2005]. A linear increase in this cost with sequestered fraction, S, is assumed, rising to 250\$ per tonne of $CO_2$ sequestered when the fraction is unity, hence $C_s$=50\$+200S\$. The costs in the region of S = 1 are highly uncertain, as they would be determined by emissions technologies such as biomass energy with CCS or direct air capture to compensate for residual emissions and mobile sources. What fraction of emissions will fall into these categories is difficult to predict. Estimates of the cost of direct air capture range from as little as 100$ per tonne to as much as ten times this amount [Kriegler *et al.,* 2013]. Whilst the technology to sequester 100% of the emissions of extracted carbon already exists, it is the cost that will determine the long-term fate of the fossil fuel industry. The technological challenge for the industry's survival is clear: if the cost of CCS is too high, fossil fuels will be replaced by more cost-effective alternatives in many applications. All IPCC scenarios, however, envisage some applications for fossil fuels for the indefinite future.

Whilst there are important uncertainties surrounding the costs near S = 1, there is greater confidence in the costs of the policy at lower sequestered fractions. In the short-to-medium term, the cost per tonne of $CO_2$-equivalent (a composite metric of multiple greenhouse gases, in which emissions of different gases are weighted by their Global Warming Potential [Shine *et al.,* 2007]) extracted, $C_e$, (reflecting the overall policy cost on the fossil fuel system; $C_e=SC_s$ as S is the fraction of the total extracted carbon which is subsequently sequestered) is given as $C_e$=S(50\$+200\$*S) and is shown in figure 7.7f. Based on this cost-curve, $C_e$ may remain relatively low over the course of the century. For instance, by the time S=0.2 in the early 2030s, $C_s$ = 90\$/t$CO_2$ and $C_e$=18\$/t$CO_2$ (see figure 7.7e). Even if the entire policy cost were borne by consumers, maintaining producer surplus, this is equivalent to an oil price increase of approximately 7-8$ per barrel, and is towards the lower end of estimates of carbon prices required to meet the 2K goal.

What might such an approach cost in total? Figure 7.7b shows that, even under the assumption that the additional cost of sequestration has no impact on fossil fuel demand and consumption of carbon follows business-as-usual (RCP8.5), the total policy cost in 2030 is 0.8$tn/yr, which is comparable to policy alternatives that also achieve the 2K goal. Exclusive reliance on CCS, without other technologies and measures reducing fossil fuel demand, would be extremely expensive later in the century, but is also entirely unrealistic: other energy technologies will be deployed in response to the increased costs of sequestration, reducing the overall policy cost.

## 7.4 Summary and conclusions

This chapter described a simple modification of a coupled climate-carbon-cycle model, in order to accurately capture the physics of the response to $CO_2$ emissions shown by ESMs. This is achieved by introducing a simple linear relationship between $iIRF_{100}$ and both GMST increase and cumulative carbon uptake in the climate system. The $iIRF_{100}$ metric provides a potential parallel to those used to assess the thermal response to radiative forcing, namely the TCR and the ECS, for assessing carbon-cycle feedbacks. The use of the 100 year time horizon for iIRF represents an inherent value choice. Whilst features of the climate response to $CO_2$ over a 100-year timescale is particularly relevant when typical economic discount rates are considered, a focus on different time horizons to summarise the response of the carbon cycle would represent other valid value-driven choices for the timescale of interest and may be more appropriate for certain applications. Although a useful composite metric for the coupled climate-carbon-cycle system exists, the TCRE, future studies of carbon-cycle behaviour could report on ranges of $\mathrm{iIRF}_{100}$, and importantly for carbon-cycle feedbacks, the evolution of this metric over time under specific emissions scenarios, in order to isolate the changing response of the carbon-cycle from uncertainty in the thermal response of the climate system to radiative forcing.

We have shown that including both explicit $CO_2$ uptake- and temperature- induced feedbacks are essential to capture ESM behaviour. Important dependences of the carbon-cycle response to pulse size, background conditions and the suppression of temperature-induced feedbacks are generally well captured with the FAIR model.

Having developed and validated the FAIR model, I use it in the remainder of the chapter to discuss a proposed adaptive mitigation policy that indexes changes in the sequestered fraction of extracted carbon, S, to changes in the human attributable warming index (AWI). S must eventually reach unity, irrespective of whether stabilisation is at 2K, 3K or 4K, and irrespective of how much renewable and nuclear energy are deployed in the meantime. This will almost certainly require the ability to capture and store carbon. S is currently close to zero and shows no sign of increasing appreciably without intervention. IPCC AR5 identified carbon capture and sequestration (CCS) as a critical technology to delivering a 2-in-3 or better chance of limiting GMST increases below 2K, categorised in IPCC [2014c] as 430-480 ppm atmospheric $CO_2$-equivalent concentration (total anthropogenic radiative forcing expressed as the concentration of $CO_2$ that would produce this same total forcing) in 2100. If CCS is excluded as a mitigation option, only 4 out of 10 of the IAMs considered in IPCC-WG3 can achieve this budget, indicating that a '2K carbon budget' may be almost infeasible without CCS (figure 6.24 of [Clarke, 2014]). As assessed by IPCC-WG3, failing to deploy CCS at scale increases the costs of meeting the '2K carbon budget' by around a factor of two, much more than excluding any other mitigation technology.

Recent work [McGlade & Ekins, 2015] suggests that deployment of CCS from 2025 has only a limited impact on allowable fossil fuel use to 2050 in scenarios that meet the 2K budget. However, this assumes a limited pace of deployment of CCS and downplays the vital role of CCS after 2050 in 2K and even 3K emission scenarios [Torvanger *et al.,* 2013], both in reducing fossil carbon emissions to zero and compensating for residual emissions through biomass energy with CCS.

Since large-scale CCS appears to be needed to meet climate goals, then climate policy should ensure that this technology is scaled up smoothly to ensure it is as safe, cost-effective and publicly acceptable as possible. This is particularly important given that the success of CCS depends on the long-term behaviour of geological storage sites [Torvanger *et al.,* 2012]. Any scenario that envisages a very rapid increase in the sequestered fraction would necessarily require the introduction of a large number of new sites simultaneously, increasing the economic, environmental and societal risks.

As an example of such a policy I propose a policy that indexes increases in S to increases in the AWI. Despite the broad uncertainty in the climate response, analysed over several previous chapters of this thesis, it remains true, independent of a particularly high or low climate response to $CO_2$ emissions, that global net emissions of long-lived climate pollutants (mainly $CO_2$) must fall to zero by the time temperatures reach the peak warming goal to be able to stabilise GMST at this level. I have shown that such an adaptive approach, where the mitigation policy responds to the emergent climate response observed in the real world, can be less hostage to uncertainty in the physical climate response with regard to limiting warming to beneath policy-relevant thresholds. An adaptive policy framework, focused on policy variables that can limit the stock of cumulative emissions in the atmosphere, could therefore offer a productive way forward for construction of mitigation policies in light of the broad and seemingly irreducible uncertainty in the physical climate response to $CO_2$.

# CHAPTER 8

## Summary and conclusions

Uncertainty about the climate response to $CO_2$ is a fundamental challenge for climate policy, in particular for constructing policies that will mitigate anthropogenic emissions of greenhouse gases to limit global-mean warming to below specified thresholds. This uncertainty arises from the complex interaction of processes across multiple spatial and temporal scales, giving rise to both amplifying and damping feedback loops (chapter 2). Therefore, as described in chapter 1, this thesis aimed to investigate and understand uncertainty in the magnitude of the physical climate response to $CO_2$, and if this uncertainty could be constrained further from currently assessed ranges.

In this thesis I have used multiple models of the climate system, simple and complex (see chapter 2), to investigate uncertainty in the physical climate response. Considering first simple models of the climate system, in chapter 4 I investigated constraints on the climate response provided by recent observations of the climate system using a simple two-box impulse-response energy-balance model. Even given identical observational data, important and often neglected sensitivities arising from structural assumptions in the energy-balance model were shown to have important impacts on estimation of the TCR uncertainty. This is particularly true for the upper bound of the uncertainty interval, which has potentially large impacts on climate policy analysis under certainty (chapter 1). Inferences about properties of the climate response from observations cannot ever be model-independent and will always be functions of the model structure used to link observations with properties of the climate response. Additionally, I showed that estimates of uncertainty in ECS using observations and energy-balance models are strongly dependent on the efficacy of ocean heat uptake, which is difficult to measure from observations alone. Future studies should focus on estimating uncertainties in the TCR (and possibly RWF) from observations, because estimates of the TCR can be better constrained through this method than for ECS and because, as showed in chapter 4, it is a more important measure of climate response to the evolution of GMST under policy-relevant future climate forcing scenarios (except in situations where the ECS is very high or significant non-linearities are encountered in the climate system response).

Chapter 4 also proposed a new metric of the climate response, the realised warming fraction or RWF (the ratio of TCR to ECS) which, together with the TCR, form approximately independent dimensions of the climate response space (unlike TCR and ECS). Climate policy analyses, such as calculations of the social cost of carbon (the additional economic damage caused by emitting an extra tonne of $CO_2$ today), often require a sampling of parameters of climate response in simple physical climate models. The use of RWF and TCR as two independent parameters of climate response (which could then be sampled independently, as done in chapter 7) could be convenient for such analyses and offers a simple way to accurately capture the joint distribution of TCR and ECS.

Describing the climate response space with RWF and TCR highlights a lack of GCMs with high RWF and low TCR (ECS<2K - see chapters 3 and 4), a region of climate response space that is consistently included within observationally-based uncertainty estimates. This absence of low climate response GCMs includes the diverse PPE from chapter 3 which contained many GCMs that have very high ($\geq$ 8K) ECS, but a lack of GCMs with ECS<2K. Whether this should be interpreted as a physical constraint on the climate response, or merely a product of insufficiently diverse GCM ensembles was not clear.

In order to explore the lower limit of climate response, chapter 5 constructed a statistical emulator of the HadCM3 parameter space consisting of a multiple linear regression model selected using an automated regressor selection algorithm. The emulator was found to have some skill in predicting low ECS GCMs that also have realistic radiative control climates. An interesting extension for this work would be to map HadCM3 parameter space using a more sophisticated statistical model, such as a 'general-to-specific' model selection technique [Pretis *et al.,* 2016]. A more complex statistical model selection algorithm might further improve understanding of important interactions between GCM parameters in determining the climate response and properties of the control climate. I find that the PPE members with ECS<2K perform surprisingly well under basic validation of their control climate simulations and would not be complete outliers from the CMIP5 MME. Additionally, unlike the higher-resolution flagship GCMs in the CMIP5 MME, these low ECS PPE members have not been tuned to optimise their present-day climate simulation whilst maintaining their ECS values. The emulator structure I developed here, or a more sophisticated statistical tool, could be used to improve the climate simulations of these GCMs without significantly affecting their ECS. Such a tuning exercise would allow a more 'like-for-like' comparison of GCM skill between the PPE members and the CMIP5 GCMs. Whilst many different parameters have already been perturbed within the HadCM3 PPE, several other parameters (discussed in more detail in chapter 5) might offer new degrees of freedom for achieving low climate sensitivity, whilst maintaining or improving control climate fidelity. These include the anvil shape factor parameter (which perturbs the cloud area due to convective anvils in the GCM radiation scheme) and the temperature dependent partitioning between liquid and ice phase condensed water, which has been shown to affect the simulated properties of SW-CRE and the cloud phase-change feedback [McCoy *et al.,* 2016]. I am currently including the ability to perturb these parameters within the climate*prediction*.net infrastructure in order to explore the effect of these parametrisations on the ECS and control cloud simulation in the HadCM3 PPE.

In chapters 5 and 6 only the pre-industrial control simulations were analysed to assess the simulation skill of the low ECS GCMs. Historical integrations, with multiple observed forcing agents, would be useful additional validation of these GCMs and in particular would reveal their sensitivity to non-$CO_2$ radiative forcing. Although it is impossible to *apriori* exclude the possibility of significantly non-unity efficacies of non-$CO_2$ forcings, as the GCM TCRs lie within the uncertainties of TCR estimated from observations over the $20^{th}$ century (chapter 4), these GCMs would be expected to approximately match the historical record (within uncertainties) for approximately unity efficacy of non-$CO_2$ forcing.

The physical mechanisms that drive the low ECS in these PPE members are complex, as described in chapter 6. The main driver of the low ECS are decreases in SW-CRE, in particular over regions of the low-latitude oceans with high atmospheric stability and populated by low clouds. The importance of cloud feedbacks in this regime is consistent with the largest sources of uncertainty in the CMIP5 MME. Low ECS is associated with too strong SW-CRE cooling over regions of low cloud cover. This is driven by too bright low clouds, although several of the low ECS GCMs have more low cloud in the control climate simulations than the standard HadCM3 and other GCMs, which simulate too little compared to observations. Low clouds are seen to grow under doubled $CO_2$. A more comprehensive diagnostic set would help to understand the mechanisms for the negative low cloud feedbacks better, as it would allow for the decomposition of the cloud feedback into different optical thickness and liquid water path bins. Diagnostics such as the

ISCCP cloud simulator [Webb *et al.,* 2001], which offers a better comparison between GCM fields and observations by applying an estimate of the satellite transfer function to the GCM output, are currently being included in these simulations to allow this analysis. Idealised experiments, in which specific process are suppressed (such as prescribed sea-surface temperature simulations, or fixed boundary layer humidity), could be conducted with the low climate response parameter sets to understand more about the physical mechanisms underlying the simulated cloud response.

In this thesis I have focused on perturbations to HadCM3 only, however, it is possible that a similar diversity of climate response could be obtained in CMIP5 GCMs if large enough perturbed ensembles could be created, that (due to the improved resolution and climate processes modelling) may have improved control climate simulations relative to the GCMs examined in this thesis. A multi-model extension of the multi-parameter multi-physics ensemble (which has been used to sample model structure through different parametrisation schemes as well as differing parameter values), recently simulated for the MIROC GCM [Shiogama *et al.,* 2014], could be a promising way to comprehensively assess intra and inter GCM uncertainty for future assessments of the climate response.

Based on the investigations into low climate response throughout this thesis, it currently seems difficult to convincingly rule out ECS in the real climate system between 1.5K and 2K, when multiple lines of evidence are assessed. Seemingly irreducible (for the moment) uncertainties exist in the magnitude of the climate response to $CO_2$, both at the high and low ends of the climate response range. Recent work [Stevens *et al.,* 2016] proposes that a strategy focused on "ruling-out" regions of climate response space, as has been attempted in this thesis, offers a potentially fruitful way forward in narrowing the uncertainty range on ECS. Whilst this thesis has focused on observationally-based and GCM modelling techniques to constrain climate response parameters, other avenues of evidence, for instance paleoclimate observations and simulations, could also be used to help rule-out regions of ECS space as inconsistent with all lines of evidence. Continued future climate change will allow the TCR and the TCRE to be constrained more tightly, particularly if emissions of short-lived forcing agents, such as tropospheric aerosols which contribute substantially to present-day total radiative forcing uncertainty, are reduced [Penner *et al.,* 2010; Myhre *et al.,* 2015] as might be expected under the RCP scenarios. However, prioritising the mitigation of short-lived species with uncertain radiative forcing at the expense of $CO_2$ mitigation in the near-term in order to more accurately constrain the climate response could risk locking-in more permanent $CO_2$-induced warming.

A clear implication of this thesis is that adapting to climate response uncertainty will be a necessary part of successful climate policy limiting warming to beneath certain thresholds. Chapter 7 concludes the thesis by offering a mitigation policy framing that can adapt to the emerging climate response as it is observed under future emissions scenarios. Simulating the consequences of adaptive policy requires the addition of a computationally inexpensive carbon-cycle model that correctly captures the feedbacks between the climate and carbon-cycle. Whilst the model constructed here (FAIR) captures the dynamics apparent in ESMs under specified concentration increase experiments, the feedbacks in the carbon-cycle under stable and falling emissions currently remain uncertain. As these feedbacks will have an important impact on adaptive mitigation policies aiming to stabilise warming at 1.5K or 2K, a better understanding of these feedbacks under falling emissions will be an important area of further work for the climate science community.

Adaptive climate scenarios arguably provide a more directly-relevant parallel to the aims of the international climate policy process and could complement the existing scenario projections framework [O'Neill *et al.,* 2016], which do not explicitly capture feedbacks between climate impacts and policy responses leading to complexities in the interpretation of these scenarios for analysis of future adaptive mitigation pathways [Rogelj *et al.,* 2016]. An important future contribution of the science community could be simulations of adaptive climate policy scenarios for various warming thresholds with a range of ESMs and GCMs from across the climate response uncertainty distribution. These simulations could help elucidate important differences in the physical climate state between a world in which temperatures have been successfully stabilised at a certain value and a world with the same global-mean warming but in which temperatures would continue to rise further in the future. Such experiments would help understand both the range of mitigation pathways required to achieve temperature stabilisation, as well as the range of climate impacts in the stabilised climate state.

# Bibliography


1. Adler, R. F., Huffman, G. J., *et al.* The version-2 global precipitation climatology project (GPCP) monthly precipitation analysis (1979-present). *Journal of Hydrometeorology* **4,** 1147–1167 (2003).
2. Akaike, H. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on* **19,** 716–723 (1974).
3. Aldrin, M., Holden, M., *et al.* Bayesian estimation of climate sensitivity based on a simple climate model fitted to observations of hemispheric temperatures and global ocean heat content. *Environmetrics* **23,** 253–271 (2012).
4. Alduchov, O. A. & Eskridge, R. E. Improved Magnus form approximation of saturation vapor pressure. *Journal of Applied Meteorology* **35,** 601–609 (1996).
5. Allen, M. Do-it-yourself climate prediction. *Nature* **401,** 642–642 (1999).
6. Allen, M. R. & Ingram, W. J. Constraints on future changes in climate and the hydrologic cycle. *Nature* **419,** 224–232 (2002).
7. Allen, M. R. & Stocker, T. F. Impact of delay in reducing carbon dioxide emissions. *Nature Climate Change* **4,** 23–26 (2014).
8. Allen, M. R., Frame, D. J. & Mason, C. F. The case for mandatory sequestration. *Nature Geoscience* **2,** 813–814 (2009).
9. Allen, M. R., Frame, D. J., *et al.* Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature* **458,** 1163–1166 (2009).
10. Andrews, D. G. & Allen, M. R. Diagnosis of climate models in terms of transient climate response and feedback response time. *Atmospheric Science Letters* **9,** 7–12 (2008).
11. Andrews, T., Gregory, J. M., Webb, M. J. & Taylor, K. E. Forcing, feedbacks and climate sensitivity in CMIP5 coupled atmosphere-ocean climate models. *Geophysical Research Letters* **39** (2012).
12. Arora, V. K., Boer, G. J., *et al.* Carbon–concentration and carbon–climate feedbacks in CMIP5 Earth system models. *Journal of Climate* **26,** 5289–5314 (2013).
13. Bindoff, N. L. *et al.* Detection and Attribution of Climate Change: from Global to Regional. *Climate Change 2013: The Physical Science Basis. Con-*

*tribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]* (2013).

14. Bloch-Johnson, J., Pierrehumbert, R. T. & Abbot, D. S. Feedback temperature dependence determines the risk of high warming. *Geophysical Research Letters* **42,** 4973–4980 (2015).
15. Bony, S. & Dufresne, J.-L. Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models. *Geophysical Research Letters* **32** (2005).
16. Bony, S., Colman, R., *et al.* How well do we understand and evaluate climate change feedback processes? *Journal of Climate* **19,** 3445–3482 (2006).
17. Boucher, O. & Reddy, M. S. Climate trade-off between black carbon and carbon dioxide emissions. *Energy Policy* **36,** 193–200 (2008).
18. Boucher, O. *et al.* Clouds and Feedback. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]* (2013).
19. Bowerman, N. H. A., Frame, D. J., *et al.* Cumulative carbon emissions, emissions floors and short-term rates of warming: implications for policy. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **369,** 45–66 (2011).
20. Bretherton, C. S. Insights into low-latitude cloud feedbacks from high-resolution models. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **373,** 20140415 (2015).
21. Brient, F. & Bony, S. How may low-cloud radiative properties simulated in the current climate influence low-cloud feedbacks under global warming? *Geophysical Research Letters* **39** (2012).
22. Brient, F. & Bony, S. Interpretation of the positive low-cloud feedback predicted by a climate model under global warming. *Climate Dynamics* **40,** 2415–2431 (2013).
23. Caldeira, K. & Kasting, J. F. Susceptibility of the early Earth to irreversible glaciation caused by carbon dioxide clouds (1992).
24. Cess, R. D., Potter, G. L., *et al.* Intercomparison and interpretation of climate feedback processes in 19 atmospheric general circulation models. *Journal of Geophysical Research* **95,** 601216 (1990).
25. Charney, J. G., Arakawa, A., *et al. Carbon dioxide and climate: a scientific assessment* 1979.
26. Church, J. A., Monselesan, D., Gregory, J. M. & Marzeion, B. Evaluating the ability of process based models to project sea-level change. *Environmental Research Letters* **8,** 014051 (2013).
27. Clarke, L. *et al.* Assessing Transformation Pathways. *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Edenhofer, O., R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J.*

*Savolainen, S. Schlmer, C. von Stechow, T. Zwickel and J.C. Minx (eds.)]* (2014).

28. Collins, M., Booth, B. B. B., *et al.* Climate model errors, feedbacks and forcings: a comparison of perturbed physics and multi-model ensembles. *Climate Dynamics* **36,** 1737–1766 (2011).
29. Collins, M., Brierley, C. M., *et al.* The sensitivity of the rate of transient climate change to ocean physics perturbations. *Journal of Climate* **20,** 2315–2320 (2007).
30. Colman, R. A. & McAvaney, B. J. On tropospheric adjustment to forcing and climate feedbacks. *Climate Dynamics* **36,** 1649–1658 (2011).
31. Cowtan, K. & Way, R. G. Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends. *Quarterly Journal of the Royal Meteorological Society* **140,** 1935–1944 (2014).
32. Dahowski, R. T., Davidson, C. L., Li, X. C. & Wei, N. A \$70/t$CO_2$ greenhouse gas mitigation backstop for China's industrial and electric power sectors: Insights from a comprehensive CCS cost curve. *International Journal of Greenhouse Gas Control* **11,** 73–85 (2012).
33. Dee, D. P., Uppala, S. M., *et al.* The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society* **137,** 553–597 (2011).
34. Dessler, A. E. A determination of the cloud feedback from climate variations over the past decade. *Science* **330,** 1523–1527 (2010).
35. Domingues, C. M., Church, J. A., *et al.* Improved estimates of upper-ocean warming and multi-decadal sea-level rise. *Nature* **453,** 1090–1093 (2008).
36. Donohoe, A., Armour, K. C., Pendergrass, A. & Battisti, D. Shortwave and longwave radiative contributions to global warming under increasing $CO_2$. *Proceedings of the National Academy of Sciences* **111,** 16700–16705 (2014).
37. Draper, N. R. & Smith, H. *Applied regression analysis* (John Wiley & Sons, 2014).
38. Dufresne, J.-L. & Bony, S. An assessment of the primary sources of spread of global warming estimates from coupled atmosphere-ocean models. *Journal of Climate* **21,** 5135–5144 (2008).
39. Edwards, P. N. *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming* (MIT Press, 2010).
40. Edwards, P. N. History of climate modeling. *Wiley Interdisciplinary Reviews: Climate Change* **2,** 128–139 (2011).
41. Engström, A., Bender, F. A.-M., Charlson, R. J. & Wood, R. The nonlinear relationship between albedo and cloud fraction on near-global, monthly mean scale in observations and in the CMIP5 model ensemble. *Geophysical Research Letters* **42,** 9571–9578 (2015).
42. Flato, G. *et al.* Evaluation of Climate Models. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]* (2013).
43. Forest, C. E., Stone, P. H., *et al.* Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science* **295,** 113–117 (2002).

44. Forster, P. M., Andrews, T., *et al.* Evaluating adjusted forcing and model spread for historical and future scenarios in the CMIP5 generation of climate models. *Journal of Geophysical Research: Atmospheres* **118,** 1139–1150 (2013).
45. Frame, D. J., Stone, D. A., Stott, P. & Allen, M. Alternatives to stabilization scenarios. *Geophysical Research Letters* **33** (2006).
46. Friedlingstein, P., Cox, P., *et al.* Climate-carbon cycle feedback analysis: Results from the C4MIP model intercomparison. *Journal of Climate* **19,** 3337–3353 (2006).
47. Fuglestvedt, J. S., Shine, K. P., *et al.* Transport impacts on atmosphere and climate: Metrics. *Atmospheric Environment* **44,** 4648–4677 (2010).
48. Geoffroy, O., Saint-Martin, D., *et al.* Transient climate response in a two-layer energy-balance model. Part II: Representation of the efficacy of deep-ocean heat uptake and validation for CMIP5 AOGCMs. *Journal of Climate* **26,** 1859–1876 (2013).
49. Gillett, N. P., Arora, V. K., Matthews, D & Allen, M. Constraining the Ratio of Global Warming to Cumulative $CO_2$ Emissions Using CMIP5 Simulations. *Journal of Climate* **26,** 6844–6858 (2013).
50. Glotter, M. J., Pierrehumbert, R. T., *et al.* A simple carbon cycle representation for economic and policy analyses. *Climatic Change* **126,** 319–335 (2014).
51. Good, P., Gregory, J. M. & Lowe, J. A step-response simple climate model to reconstruct and interpret AOGCM projections. *Geophysical Research Letters* **38** (2011).
52. Good, P., Gregory, J. M., Lowe, J. & Andrews, T. Abrupt $CO_2$ experiments as tools for predicting and understanding CMIP5 representative concentration pathway projections. *Climate Dynamics* **40,** 1041–1053 (2013).
53. Gordon, C., Cooper, C., *et al.* The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre coupled model without flux adjustments. *Climate Dynamics* **16,** 147–168 (2000).
54. Gregory, D. & Rowntree, P. R. A mass flux convection scheme with representation of cloud ensemble characteristics and stability-dependent closure. *Monthly Weather Review* **118,** 1483–1506 (1990).
55. Gregory, D., Inness, P. & Gregory, J. M. *Unified Model Documentation Paper 26: Large-scale Precipitation* tech. rep. (Met Office, 1996).
56. Gregory, D., Inness, P. & Gregory, J. M. *Unified Model Documentation Paper 27: Convection Scheme* tech. rep. (Met Office, 1996).
57. Gregory, J. M. & Forster, P. M. Transient climate response estimated from radiative forcing and observed temperature change. *Journal of Geophysical Research: Atmospheres (1984–2012)* **113** (2008).
58. Gregory, J. M. & Webb, M. J. Tropospheric adjustment induces a cloud component in $CO_2$ forcing. *Journal of Climate* **21,** 58–71 (2008).
59. Gregory, J. M., Ingram, W. J., *et al.* A new method for diagnosing radiative forcing and climate sensitivity. *Geophysical Research Letters* **31** (2004).
60. Gregory, J. M., Stouffer, R. J., *et al.* An observationally based estimate of the climate sensitivity. *Journal of Climate* **15,** 3117–3121 (2002).
61. Gregory, J. M., Jones, C. D., Cadule, P & Friedlingstein, P. Quantifying carbon cycle feedbacks. *Journal of Climate* **22,** 5232–5250 (2009).

62. Gregory, J. M., Andrews, T. & Good, P. The inconstancy of the transient climate response parameter under increasing $CO_2$. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **373,** 20140417 (2015).
63. Hadley, G. Concerning the cause of the general trade winds. *Philosophical Transactions of the Royal Society of London* **39,** 58–62 (1735).
64. Halley, E. An historical account of the trade winds, and monsoons, observable in the seas between and near the tropicks, with an attempt to assign the physical cause of said winds. *Transactions of the Royal Society of London* **16,** 153–168 (1686).
65. Hansen, J. *et al.* in *Climate Processes and Climate Sensitivity* (American Geophysical Union, 1984).
66. Hansen, J., Sato, M., *et al.* Efficacy of climate forcings. *Journal of Geophysical Research: Atmospheres (1984–2012)* **110** (2005).
67. Hansen, J., Sato, M. & Ruedy, R. Radiative forcing and climate response. *Journal of Geophysical Research: Atmospheres (1984–2012)* **102,** 6831–6864 (1997).
68. Harris, G. R., Sexton, D. M. H., *et al.* Probabilistic projections of transient climate change. *Climate Dynamics* **40,** 2937–2972 (2013).
69. Hartmann, D. L. & Larson, K. An important constraint on tropical cloud-climate feedback. *Geophysical Research Letters* **29,** 12–1 (2002).
70. Hartmann, D. L. & Michelsen, M. L. No evidence for iris. *Bulletin of the American Meteorological Society* **83,** 249–254 (2002).
71. Hartmann, D. L., Moy, L. A. & Fu, Q. Tropical convection and the energy balance at the top of the atmosphere. *Journal of Climate* **14,** 4495–4511 (2001).
72. Hawkins, E. & Sutton, R. The potential to narrow uncertainty in regional climate predictions. *Bulletin of the American Meteorological Society* **90,** 1095–1107 (2009).
73. Hegerl, G. C. *et al.* Understanding and Attributing Climate Change. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]* (2007).
74. Held, I. M., Winton, M., *et al.* Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing. *Journal of Climate* **23,** 2418–2427 (2010).
75. Herrington, T. & Zickfeld, K. Path dependence of climate and carbon cycle response over a broad range of cumulative carbon emissions. *Earth System Dynamics Discussion* **5,** 747–778 (2014).
76. Hof, A. F., Hope, C. W., *et al.* The benefits of climate change mitigation in integrated assessment models: the role of the carbon cycle and climate component. *Climatic Change* **113,** 897–917 (2012).
77. Huber, M. & Knutti, R. Natural variability, radiative forcing and climate response in the recent hiatus reconciled. *Nature Geoscience* (2014).
78. Huntingford, C. Climate projection: refining global warming projections. *Nature Climate Change* **3,** 704–705 (2013).

79. Huntingford, C. & Cox, P. M. An analogue model to derive additional climate change scenarios from existing GCM simulations. *Climate Dynamics* **16,** 575–586 (2000).
80. Huntingford, C., Lowe, J. A., *et al.* Contributions of carbon cycle uncertainty to future climate projection spread. *Tellus B* **61,** 355–360 (2009).
81. Ingram, W. J. A very simple model for the water vapour feedback on climate change. *Quarterly Journal of the Royal Meteorological Society* **136,** 30–40 (2010).
82. Ingram, W. J. Some implications of a new approach to the water vapour feedback. *Climate Dynamics* **40,** 925–933 (2013).
83. IPCC. Summary for Policymakers. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Forth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]* (2007).
84. IPCC. Summary for Policymakers. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]* (2013).
85. IPCC. Summary for Policymakers. *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Field, C.B., V.R. Barros, D.J. Dokken, K.J. Mach, M.D. Mastrandrea, T.E. Bilir, M. Chatterjee, K.L. Ebi, Y.O. Estrada, R.C. Genova, B. Girma, E.S. Kissel, A.N. Levy, S. MacCracken, P.R. Mastrandrea, and L.L. White (eds.)]* (2014).
86. IPCC. Summary for Policymakers. *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Edenhofer, O., R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlmer, C. von Stechow, T. Zwickel and J.C. Minx (eds.)]* (2014).
87. IPCC. Summary for Policymakers. *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)]* (2014).
88. Jackson, L. C., Vellinga, M. & Harris, G. R. The sensitivity of the meridional overturning circulation to modelling uncertainty in a perturbed physics ensemble without flux adjustment. *Climate Dynamics* **39,** 277–285 (2012).
89. Jensen, S. & Traeger, C. P. in *Working Paper* (2013).
90. Jones, G. S., Stott, P. A. & Christidis, N. Attribution of observed historical near–surface temperature variations to anthropogenic and natural causes using CMIP5 simulations. *Journal of Geophysical Research: Atmospheres* **118,** 4001–4024 (2013).
91. Joos, F., Bruno, M., *et al.* An efficient and accurate representation of complex oceanic and biospheric models of anthropogenic carbon uptake. *Tellus B* **48,** 397–417 (1996).

92. Joos, F., Roth, R., *et al.* Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmospheric Chemistry and Physics* **13,** 2793–2825 (2013).
93. Kamae, Y. & Watanabe, M. On the robustness of tropospheric adjustment in CMIP5 models. *Geophysical Research Letters* **39** (2012).
94. Klein, S. A. & Hall, A. Emergent constraints for cloud feedbacks. *Current Climate Change Reports* **1,** 276–287 (2015).
95. Klein, S. A. & Hartmann, D. L. The seasonal cycle of low stratiform clouds. *Journal of Climate* **6,** 1587–1606 (1993).
96. Knutti, R. The end of model democracy? *Climatic Change* **102,** 395–404 (2010).
97. Knutti, R. & Rogelj, J. The legacy of our $CO_2$ emissions: a clash of scientific facts, politics and ethics. *Climatic Change* **133,** 361–373 (2015).
98. Knutti, R., Furrer, R., *et al.* Challenges in combining projections from multiple climate models. *Journal of Climate* **23,** 2739–2758 (2010).
99. Knutti, R., Krähenmann, S., Frame, D. & Allen, M. Comment on "Heat capacity, time constant, and sensitivity of Earth's climate system" by SE Schwartz. *Journal of Geophysical Research: Atmospheres (1984–2012)* **113** (2008).
100. Knutti, R., Joos, F., *et al.* Probabilistic climate change projections for $CO_2$ stabilization profiles. *Geophysical Research Letters* **32** (2005).
101. Kravitz, B., Robock, A., *et al.* The geoengineering model intercomparison project (GeoMIP). *Atmospheric Science Letters* **12,** 162–167 (2011).
102. Kriegler, E., Edenhofer, O., *et al.* Is atmospheric carbon dioxide removal a game changer for climate change mitigation? *Climatic Change* **118,** 45–57 (2013).
103. Kummer, J. R. & Dessler, A. E. The impact of forcing efficacy on the equilibrium climate sensitivity. *Geophysical Research Letters* **41,** 3565–3568 (2014).
104. Larson, K., Hartmann, D. L. & Klein, S. A. The role of clouds, water vapor, circulation, and boundary layer structure in the sensitivity of the tropical climate. *Journal of Climate* **12,** 2359–2374 (1999).
105. Le Quéré, C., Moriarty, R., *et al.* Global Carbon Budget 2015. *Earth System Science Data* **7,** 349–396 (2015).
106. Leduc, M., Matthews, H. D. & de Elía, R. Quantifying the limits of a linear temperature response to cumulative $CO_2$ emissions. *Journal of Climate* **28,** 9955–9968 (2015).
107. Levitus, S., Antonov, J. I., *et al.* World ocean heat content and thermosteric sea level change (0–2000 m), 1955–2010. *Geophysical Research Letters* **39** (2012).
108. Lewis, N. An Objective Bayesian Improved Approach for Applying Optimal Fingerprint Techniques to Estimate Climate Sensitivity. *Journal of Climate* **26** (2013).
109. Lewis, N. & Curry, J. A. The implications for climate sensitivity of AR5 forcing and heat uptake estimates. *Climate Dynamics,* 1–15 (2014).
110. Li, S. & Jarvis, A. Long run surface temperature dynamics of an A-OGCM: the HadCM3 4× $CO_2$ forcing experiment revisited. *Climate Dynamics* **33,** 817–825 (2009).

111. Lindzen, R. S., Chou, M.-D. & Hou, A. Y. Does the earth have an adaptive infrared iris? *Bulletin of the American Meteorological Society* **82,** 417–432 (2001).
112. Loeb, N. G., Wielicki, B. A., *et al.* Toward optimal closure of the Earth's top-of-atmosphere radiation budget. *Journal of Climate* **22,** 748–766 (2009).
113. Lunt, D. J., Haywood, A. M., *et al.* Earth system sensitivity inferred from Pliocene modelling and data. *Nature Geoscience* **3,** 60–64 (2010).
114. Manabe, S. & Strickler, R. F. Thermal equilibrium of the atmosphere with a convective adjustment. *Journal of the Atmospheric Sciences* **21,** 361–385 (1964).
115. Manabe, S. & Wetherald, R. T. Thermal Equilibrium of the Atmosphere with a Given Distribution of Relative Humidity. *Journal of the Atmospheric Sciences* (1967).
116. Matthews, H. D. Quantifying historical carbon and climate debts among nations. *Nature Climate Change* (2015).
117. Matthews, H. D. & Caldeira, K. Stabilizing climate requires near-zero emissions. *Geophysical Research Letters* **35** (2008).
118. Mauritsen, T. & Stevens, B. Missing iris effect as a possible cause of muted hydrological change and high climate sensitivity in models. *Nature Geoscience* **8,** 346–351 (2015).
119. Mauritsen, T., Stevens, B., *et al.* Tuning the climate of a global model. *Journal of Advances in Modeling Earth Systems* **4** (2012).
120. McCoy, D. T., Tan, I., *et al.* On the relationships among cloud cover, mixed-phase partitioning, and planetary albedo in GCMs. *Journal of Advances in Modeling Earth Systems* (2016).
121. McGlade, C. & Ekins, P. The geographical distribution of fossil fuels unused when limiting global warming to 2°C. *Nature* **517,** 187–190 (2015).
122. McKay, M. D., Beckman, R. J. & Conover, W. J. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **42,** 55–61 (2000).
123. Meehl, G. A., Covey, C., *et al.* The WCRP CMIP3 multimodel dataset: A new era in climate change research. *Bulletin of the American Meteorological Society* **88,** 1383–1394 (2007).
124. Meinshausen, M., Meinshausen, N., *et al.* Greenhouse-gas emission targets for limiting global warming to 2°C. *Nature* **458,** 1158–1162 (2009).
125. Meinshausen, M., Smith, S. J., *et al.* The RCP greenhouse gas concentrations and their extensions from 1765 to 2300. *Climatic Change* **109,** 213–241 (2011).
126. Metz, B., Davidson, O., *et al.* IPCC Special Report on Carbon Capture and Storage. *Intergovernmental Panel on Climate Change, Working Group III* (2005).
127. Millar, R., Allen, M., Rogelj, J. & Friedlingstein, P. The cumulative carbon budget and its implications. *Oxford Review of Economic Policy* **32,** 323–342 (2016).
128. Miller, R. L. Tropical thermostats and low cloud cover. *Journal of Climate* **10,** 409–440 (1997).
129. Mitchell, J. F. B., Johns, T. C., *et al.* Towards the construction of climate change scenarios. *Climatic Change* **41,** 547–581 (1999).

130. Morice, C. P., Kennedy, J. J., Rayner, N. & Jones, P. Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set. *Journal of Geophysical Research: Atmospheres (1984–2012)* **117** (2012).
131. Murphy, J. M., Sexton, D. M. H., *et al.* Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature* **430,** 768–772 (2004).
132. Murphy, J. M., Sexton, D. M. H., *et al.* UK climate projections science report: climate change projections (2009).
133. Myhre, G. *et al.* Anthropogenic and Natural Radiative Forcing. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]* (2013).
134. Myhre, G. *et al.* Anthropogenic and Natural Radiative Forcing: Supplementary Material. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]* (2013).
135. Myhre, G., Boucher, O., *et al.* Declining uncertainty in transient climate response as $CO_2$ forcing dominates future climate change. *Nature Geoscience* **8,** 181–185 (2015).
136. Nam, C., Bony, S., Dufresne, J.-L. & Chepfer, H. The 'too few, too bright' tropical low-cloud problem in CMIP5 models. *Geophysical Research Letters* **39** (2012).
137. Nordhaus, W. D. *A question of balance: Weighing the options on global warming policies* (Yale University Press, 2008).
138. O'Neill, B. C., Tebaldi, C., *et al.* The Scenario Model Intercomparison Project (ScenarioMIP) for CMIP6. *Geoscientific Model Development Discussions* **2016,** 1–35 (2016).
139. Otto, A., Otto, F. E. L., *et al.* Energy budget constraints on climate response. *Nature Geoscience* **6,** 415–416 (2013).
140. Otto, F. E. L., Frame, D. J., Otto, A & Allen, M. Embracing uncertainty in climate change policy. *Nature Climate Change* (2015).
141. Padilla, L. E., Vallis, G. K. & Rowley, C. Probabilistic Estimates of Transient Climate Sensitivity Subject to Uncertainty in Forcing and Natural Variability. *Journal of Climate* **24** (2011).
142. Penner, J. E., Prather, M. J., *et al.* Short-lived uncertainty? *Nature Geoscience* **3,** 587–588 (2010).
143. Pierrehumbert, R. T. *Principles of Planetary Climate* (Cambridge University Press, 2010).
144. Pierrehumbert, R. T. Short-lived climate pollution. *Annual Review of Earth and Planetary Sciences* **42,** 341 (2014).
145. Pierrehumbert, R. T. Thermostats, radiator fins, and the local runaway greenhouse. *Journal of the Atmospheric Sciences* **52,** 1784–1806 (1995).
146. Pierrehumbert, R. T., Brogniez, H. & Roca, R. *On the relative humidity of the Earths atmosphere* 2007.

147. Pretis, F., Reade, J. & Sucarrat, G. *General-to-Specific (GETS) Modelling And Indicator Saturation With The R Package Gets* tech. rep. (Department of Economics Discussion Paper Series, University of Oxford., 2016).
148. Pueyo, S. Solution to the paradox of climate sensitivity. *Climatic Change* **113,** 163–179 (2012).
149. Randell, D. A. *et al.* Climate Models and Their Evaluation. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Forth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]* (2007).
150. Raper, S. C. B., Gregory, J. M. & Stouffer, R. J. The role of climate sensitivity and ocean heat uptake on AOGCM transient temperature response. *Journal of Climate* **15,** 124–130 (2002).
151. Rayner, N. A., Parker, D. E., *et al.* Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century. *Journal of Geophysical Research: Atmospheres (1984–2012)* **108** (2003).
152. Revelle, R. & Suess, H. E. Carbon dioxide exchange between atmosphere and ocean and the question of an increase of atmospheric $CO_2$ during the past decades. *Tellus* **9,** 18–27 (1957).
153. Riahi, K., Rao, S., *et al.* RCP 8.5A scenario of comparatively high greenhouse gas emissions. *Climatic Change* **109,** 33–57 (2011).
154. Richardson, M., Cowtan, K., Hawkins, E. & Stolpe, M. B. Reconciled climate response estimates from climate models and the energy budget of Earth. *Nature Climate Change* (2016).
155. Ricke, K. L. & Caldeira, K. Maximum warming occurs about one decade after a carbon dioxide emission. *Environmental Research Letters* **9,** 124002 (2014).
156. Ricke, K. L., Moreno-Cruz, J. B., *et al.* Policy thresholds in mitigation. *Nature Geoscience* **9,** 5–6 (2016).
157. Ringer, M. A., Andrews, T. & Webb, M. Global-mean radiative feedbacks and forcing in atmosphere-only and coupled atmosphere-ocean climate change experiments. *Geophysical Research Letters* **41,** 4035–4042 (2014).
158. Rogelj, J., Schaeffer, M., *et al.* Differences between carbon budget estimates unravelled. *Nature Climate Change* **6,** 245–252 (2016).
159. Rogelj, J., Hare, W., Chen, C & Meinshausen, M. Discrepancies in historical emissions point to a wider 2020 gap between 2°C benchmarks and aggregated national mitigation pledges. *Environmental Research Letters* **6,** 024002 (2011).
160. Rogelj, J., Meinshausen, M. & Knutti, R. Global warming under old and new scenarios using IPCC climate sensitivity range estimates. *Nature Climate Change* **2,** 248–253 (2012).
161. Rougier, J. Probabilistic inference for future climate using an ensemble of climate model evaluations. *Climatic Change* **81,** 247–264 (2007).
162. Rougier, J., Sexton, D. M. H., Murphy, J. & Stainforth, D. Analyzing the climate sensitivity of the HadSM3 climate model using ensembles from different but related experiments. *Journal of Climate* **22,** 3540–3557 (2009).

163. Rowlands, D. J., Frame, D. J., *et al.* Broad range of 2050 warming from an observationally constrained large climate model ensemble. *Nature Geoscience* **5,** 256–260 (2012).
164. Sanderson, B. M., Piani, C., *et al.* Towards constraining climate sensitivity by linear analysis of feedback patterns in thousands of perturbed-physics GCM simulations. *Climate Dynamics* **30,** 175–190 (2008).
165. Schaller, N., Sedláček, J. & Knutti, R. The asymmetry of the climate system's response to solar forcing changes and its implications for geoengineering scenarios. *Journal of Geophysical Research: Atmospheres* **119** (2014).
166. Schmidt, G. A., Shindell, D. T. & Tsigaridis, K. Reconciling warming trends. *Nature Geoscience* **7,** 158–160 (2014).
167. Schwartz, S. E. Heat capacity, time constant, and sensitivity of Earth's climate system. *Journal of Geophysical Research: Atmospheres (1984–2012)* **112** (2007).
168. Senior, C. A. & Mitchell, J. F. B. The time-dependence of climate sensitivity. *Geophysical Research Letters* **27,** 2685–2688 (2000).
169. Sexton, D. M. H., Murphy, J. M., Collins, M & Webb, M. Multivariate probabilistic projections using imperfect climate models part I: outline of methodology. *Climate Dynamics* **38,** 2513–2542 (2012).
170. Sherwood, S. C., Bony, S., *et al.* Adjustments in the forcing-feedback framework for understanding climate change. *Bulletin of the American Meteorological Society* (2014).
171. Sherwood, S. C., Bony, S. & Dufresne, J.-L. Spread in model climate sensitivity traced to atmospheric convective mixing. *Nature* **505,** 37–42 (2014).
172. Sherwood, S. C., Roca, R., Weckwerth, T. & Andronova, N. Tropospheric water vapor, convection, and climate. *Reviews of Geophysics* **48** (2010).
173. Shindell, D., Kuylenstierna, J. C. I., *et al.* Simultaneously mitigating near-term climate change and improving human health and food security. *Science* **335,** 183–189 (2012).
174. Shindell, D. T. Inhomogeneous forcing and transient climate sensitivity. *Nature Climate Change* **4,** 274–277 (2014).
175. Shindell, D. T., Lamarque, J.-F., *et al.* Radiative forcing in the ACCMIP historical and future climate simulations. *Atmospheric Chemistry and Physics* **13,** 2939–2974 (2013).
176. Shine, K. P., Cook, J., Highwood, E. & Joshi, M. An alternative to radiative forcing for estimating the relative importance of climate change mechanisms. *Geophysical Research Letters* **30** (2003).
177. Shine, K. P., Berntsen, T. K., *et al.* Comparing the climate effect of emissions of short-and long-lived climate agents. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **365,** 1903–1914 (2007).
178. Shiogama, H., Watanabe, M., *et al.* Multi-parameter multi-physics ensemble (MPMPE): a new approach exploring the uncertainties of climate sensitivity. *Atmospheric Science Letters* **15,** 97–102 (2014).
179. Skeie, R. B., Berntsen, T., *et al.* A lower and more constrained estimate of climate sensitivity using updated observations and detailed radiative forcing time series. *Earth System Dynamics* **5,** 139 (2014).

180. Smith, S. M., Lowe, J. A., *et al.* Equivalence of greenhouse-gas emissions for peak temperature limits. *Nature Climate Change* **2,** 535–538 (2012).
181. Soden, B. J. & Held, I. M. An assessment of climate feedbacks in coupled ocean-atmosphere models. *Journal of Climate* **19,** 3354–3360 (2006).
182. Solomon, S., Plattner, G.-K., Knutti, R. & Friedlingstein, P. Irreversible climate change due to carbon dioxide emissions. *Proceedings of the National Academy of Sciences* **106,** 1704–1709 (2009).
183. Solomon, S. *et al.* Technical Summary. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]* (2007).
184. Stainforth, D. A., Aina, T., *et al.* Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature* **433,** 403–406 (2005).
185. Stephens, G. L. Cloud feedbacks in the climate system: A critical review. *Journal of Climate* **18,** 237–273 (2005).
186. Stevens, B., Sherwood, S. C., Bony, S. & Webb, M. J. Prospects for narrowing bounds on earths equilibrium climate sensitivity. *Earth's Future* (2016).
187. Taylor, K. E. Summarizing multiple aspects of model performance in a single diagram. *Journal of Geophysical Research: Atmospheres* **106,** 7183–7192 (2001).
188. Taylor, K. E., Stouffer, R. J. & Meehl, G. A. An overview of CMIP5 and the experiment design. *Bulletin of the American Meteorological Society* **93,** 485–498 (2012).
189. Tebaldi, C. & Knutti, R. The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **365,** 2053–2075 (2007).
190. Tebaldi, C., Smith, R. L., Nychka, D. & Mearns, L. O. Quantifying uncertainty in projections of regional climate change: A Bayesian approach to the analysis of multimodel ensembles. *Journal of Climate* **18,** 1524–1540 (2005).
191. Tett, S. F. B., Rowlands, D. J., Mineter, M. J. & Cartis, C. Can top-of-atmosphere radiation measurements constrain climate predictions? Part II: Climate sensitivity. *Journal of Climate* **26,** 9367–9383 (2013).
192. Tian, B. Spread of model climate sensitivity linked to double-Intertropical Convergence Zone bias. *Geophysical Research Letters* (2015).
193. Torvanger, A., Lund, M. T. & Rive, N. Carbon capture and storage deployment rates: needs and feasibility. *Mitigation and Adaptation Strategies for Global Change* **18,** 187–205 (2013).
194. Torvanger, A., Grimstad, A.-A., *et al.* Quality of geological $CO_2$ storage to avoid jeopardizing climate targets. *Climatic Change* **114,** 245–260 (2012).
195. UNFCCC. *Paris Agreement* 2015.
196. Van Vuuren, D. P. *et al.* The representative concentration pathways: an overview. *Climatic Change* **109,** 5–31 (2011).
197. Vellinga, M. & Wu, P. Relations between northward ocean and atmosphere energy transports in a coupled climate model. *Journal of Climate* **21,** 561–575 (2008).

198. Vial, J., Dufresne, J.-L. & Bony, S. On the interpretation of inter-model spread in CMIP5 climate sensitivity estimates. *Climate Dynamics* **41,** 3339–3362 (2013).
199. Victor, D. G. & Kennel, C. F. Climate policy: Ditch the 2°C warming goal. *Nature* **514,** 30 (2014).
200. Weaver, A. J., Eby, M., *et al.* The UVic Earth System Climate Model: Model description, climatology, and applications to past, present and future climates. *Atmosphere-Ocean* **39,** 361–428 (2001).
201. Webb, M. J. & Lock, A. P. Coupling between subtropical cloud feedback and the local hydrological cycle in a climate model. *Climate Dynamics* **41,** 1923–1939 (2013).
202. Webb, M. J., Senior, C., Bony, S. & Morcrette, J.-J. Combining ERBE and ISCCP data to assess clouds in the Hadley Centre, ECMWF and LMD atmospheric climate models. *Climate Dynamics* **17,** 905–922 (2001).
203. Webb, M. J., Senior, C. A., *et al.* On the contribution of local feedback mechanisms to the range of climate sensitivity in two GCM ensembles. *Climate Dynamics* **27,** 17–38 (2006).
204. Webb, M. J., Lambert, F. H. & Gregory, J. M. Origins of differences in climate sensitivity, forcing and feedback in climate models. *Climate Dynamics* **40,** 677–707 (2013).
205. Webb, M. J., Lock, A. P., *et al.* The impact of parametrized convection on cloud feedback. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **373,** 20140414 (2015).
206. Weitzman, M. L. Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy* **5,** 275–292 (2011).
207. Wetherald, R. T. & Manabe, S. Cloud cover and climate sensitivity. *Journal of the Atmospheric Sciences* **37,** 1485–1510 (1980).
208. Williams, R. G. & Follows, M. J. *Ocean dynamics and the carbon cycle: Principles and mechanisms* (Cambridge University Press, 2011).
209. Williamson, D., Goldstein, M., *et al.* History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. *Climate Dynamics* **41,** 1703–1729 (2013).
210. Winton, M., Takahashi, K. & Held, I. M. Importance of ocean heat uptake efficacy to transient climate change. *Journal of Climate* **23,** 2333–2344 (2010).
211. Wood, R. & Bretherton, C. S. On the relationship between stratiform low cloud cover and lower-tropospheric stability. *Journal of Climate* **19,** 6425–6432 (2006).
212. Wyant, M. C., Bretherton, C. S. & Blossey, P. N. Subtropical low cloud response to a warmer climate in a superparameterized climate model. Part I: Regime sorting and physical mechanisms. *Journal of Advances in Modeling Earth Systems* **1** (2009).
213. Yamazaki, K., Rowlands, D. J., *et al.* Obtaining diverse behaviors in a climate model without the use of flux adjustments. *Journal of Geophysical Research: Atmospheres* **118,** 2781–2793 (2013).

214. Yokohata, T., Emori, S., *et al.* Climate response to volcanic forcing: Validation of climate sensitivity of a coupled atmosphere-ocean general circulation model. *Geophysical Research Letters* **32** (2005).
215. Yokohata, T., Webb, M. J., *et al.* Structural similarities and differences in climate responses to $CO_2$ increase between two perturbed physics ensembles. *Journal of Climate* **23,** 1392–1410 (2010).
216. Zelinka, M. D. & Hartmann, D. L. Why is longwave cloud feedback positive? *Journal of Geophysical Research: Atmospheres (1984–2012)* **115** (2010).
217. Zhao, M., Golaz, J.-C, *et al.* Uncertainty in model climate sensitivity traced to representations of cumulus precipitation microphysics. *Journal of Climate* **29,** 543–560 (2016).
218. Zickfeld, K. & Herrington, T. The time lag between a carbon dioxide emission and maximum warming increases with the size of the emission. *Environmental Research Letters* **10,** 031001 (2015).