Emission Targets For Avoiding Dangerous Climate Change



Niel Bowerman

Linacre College

Atmospheric, Oceanic and Planetary Physics

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2013

Emission Targets For Avoiding Dangerous Climate Change

Niel Bowerman, Linacre College, University of Oxford
Submitted for the degree of Doctor of Philosophy, Michaelmas 2013

A number of recent studies have found a strong link between peak global warming due to anthropogenic carbon dioxide and cumulative carbon emissions from the start of the industrial revolution. This thesis builds on this work by using a simple climate model to apply the concept of cumulative emissions to emission floors, by comparing cumulative emissions with other types of emissions target, and by extending the work to apply to non-carbon dioxide (CO_2) greenhouse gases and short-lived climate forcers (SLCFs).

Though peak global warming correlates well with cumulative carbon emissions, the link to emissions over shorter periods or in the years 2020 or 2050 is shown to be weaker. It is also shown that the introduction of emissions floors does not reduce the importance of cumulative emissions, but may make some warming targets unachievable. For pathways that give a most likely warming up to about 4°C, cumulative emissions from pre-industrial times to year 2200 correlate strongly with most likely resultant peak warming in the simple model used, regardless of the type of emissions floor used. The maximum rate of CO_2-induced warming is not determined by cumulative emissions but is shown to be limited by the peak rate of CO_2 emissions. A simple model of non-CO_2 greenhouse gases is also developed and used to investigate SLCFs. It is shown that emissions of SLCFs today have little impact on peak warming, and that delaying near-term reductions in SLCFs would not have a significant impact on peak warming. Only once CO_2 emissions are falling do SLCF emissions have a significant impact on peak warming. A global climate policy framework is presented as an example of how the work in this thesis could be used in policy. Future work is also discussed, particularly verification of these results in a more complex model.

Acknowledgements

This thesis was funded by a Natural Environment Research Council (NERC) CASE studentship with the Met Office. Chapters 1, 2, 4 and 5 of this report were supported in part by AVOID, a research programme funded by the UK Department of Energy and Climate Change (DECC) and the UK Department for Environment, Food and Rural Affairs (DEFRA).

I'd like to thank Myles Allen and Dave Frame, aka Framo, for their unfaltering support and belief in me over the years, and for persuading me that climate change isn't the most pressing problem in the world.

Thanks also go out to my various collaborators. To Jason Lowe for the many insightful ideas and introductions to fascinating people like Steve Smith, and of course for stumping up the cash so that I can pay my bills over the past three years. To Chris Huntingford for our many long and varied conversations, fruitful collaborations, and numerous hot chocolates. To several experts and collaborators for our influential conversations: Drew Shindell, Damon Matthews, Laila Gohar, Corinne Le Quéré, and of course to my examiners Laura Zanna and Piers Forster for their many useful recommendations. To James Annan for taking me to Japan to experience the fifth largest earthquake the world has felt in over a century, for helping Hugo Lambert and me survive that earthquake, and to both James and Hugo for bursting into laughter live on air in the background of my Radio 5 Live interview.

Finally, thanks to: Myles Allen and Dave Frame for creating the box-diffusion model; Don Grainger for creating Figure 6.4; Ed Hawkins for creating the plot that Figures 3.3 and 6.7 were adapted from; Betsy Carpenter, Su Bowerman and Amber Morgan for proof reading this tome; Neil Massey for the Linux help; Ben Bronselaer for explaining the Revelle buffer factor and for letting me constantly barge into his office to run ideas by him; William Ingram for giving me open access to his encyclopaedic knowledge; and to Dan Rowlands, Jara Imbers, Simon Driscoll, Adam Levy, Richard Millar and the other changing residents of the MAD room for keeping me entertained, from rock climbing lessons on tables in the office, to climbing into Uni Parks when we needed a breather; to Will Crouch, Michelle Hutchinson, Ben Todd, Holly Morgan, Toby Ord, Rob Wiblin, Tom Ash, and the whole team at the Centre for Effective Altruism for allowing me to design my dream job for when I graduate as an incentive to finish this thesis; to my family for helping me get this far in the first place; and finally to Elly James for keeping me sane throughout the thesis-writing project, for the trips away to look forward to, for the wild swims, and for your never-ending love and support.

Disclaimer

The views expressed herein represent those of the author and do not necessarily represent the views of the Department for Energy and Climate Change (DECC), the Department for Food, Agriculture and Rural Affairs (DEFRA), the Met Office, the Natural Environment Research Council (NERC) or the University of Oxford.

Contents

1 Introduction **8**

1.1 Uncertainty in projections of global mean warming 8

1.2 Cumulative carbon emissions 11

1.3 Climate policy 16

1.4 Modelling global warming 21

1.5 Simple models 25

1.6 Thesis outline 30

1.7 Relation to published work 32

2 Modelling the coupled carbon cycle-climate system **34**

2.1 Carbon cycle 34

2.2 The Bern simple carbon cycle model 36

2.3 A simple box-diffusion carbon cycle model 37

2.4 A simple energy balance model 54

2.5 Carbon emissions 65

2.6 Uncertainty 69

2.7 Ongoing and future work 77

2.8 Chapter summary 79

3 Testing the box-diffusion model **83**

3.1 Testing climate models 83

3.2 Testing against historical data 86

3.3 Comparison with GCMs 92
3.4 Comparison with EMICs for a pulse emission 95
3.5 Metrics 99
3.6 Symmetric and asymmetric carbon cycle feedbacks 109
3.7 Limits and applicability of the model 111

4 Emissions targets **119**

4.1 Motivation 119
4.2 Comparing emissions metrics 120
4.3 Extreme emissions pathways 125
4.4 The curvature of cumulative emissions vs peak warming 128
4.5 Emissions targets to constrain the rate of warming 130
4.6 Ongoing and future work 134
4.7 Discussion 137

5 Emissions floors **139**

5.1 Introducing emissions floors 139
5.2 Emission floors and cumulative emissions 141
5.3 Likelihoods 147
5.4 Ongoing and future work 149
5.5 Discussion and conclusions 151

6 Modelling non-CO_2 forcing agents **153**

6.1 Modelling concentrations 153
6.2 Modelling radiative forcing 163
6.3 Modelling black carbon 169
6.4 Uncertainty 173
6.5 Testing our treatment of non-CO_2 greenhouse gases 174

7 Short-lived climate forcers **179**

7.1 Short- and long-lived climate forcers 179

7.2 Action on short-lived climate forcers 181

7.3 Methods 182

7.4 Delaying SLCF reductions 183

7.5 How much do short-lived climate forcers matter for peak warming? 184

7.6 The role of short-lived climate forcers in meeting temperature goals 191

7.7 Ongoing and future work 198

7.8 Conclusions 199

8 Conclusions **201**

8.1 Thesis aims 201

8.2 Estimating the transient climate response to emissions 202

8.3 Designing a policy based on cumulative emissions 204

8.4 Comparison with other work 215

8.5 Policy implications 217

8.6 Conclusion 218

References **219**

Chapter 1

Introduction

1.1 Uncertainty in projections of global mean warming

Many things in this world are uncertain. Much of science is devoted to assessing and reducing these uncertainties, and the field of climate prediction is no different. There are many types of uncertainty in climate prediction (Cox & Stephenson 2007; Hawkins & Sutton 2009), and these can be thought of in five broad categories:

- *Initial condition uncertainty* arises from a lack of knowledge about the current state of the weather and climate. On shorter timescales, this is the dominant source of uncertainty. This source of uncertainty arises from the chaotic nature of the weather system, and the fact that initial condition errors grow exponentially on short timescales (Teixeira *et al.* 2007; Magnusson & Källén 2013).

- *Process uncertainty* arises from our inability to represent all of the components of the processes that drive climate in our models. This type of uncertainty is the most difficult to quantify, because we are unable to determine the impact of processes that are missing from a model without including them in the model in the first place (Stocker 2011).

- *Parameter uncertainty* is due to our lack of knowledge of the true values of the parameters in our models. Process uncertainty and parameter uncertainty are sometimes combined into a single "response uncertainty" or "model uncertainty".

- *Scenario uncertainty* arises because we cannot accurately predict anthropogenic emissions of greenhouse gases into the future (Le Quéré *et al.* 2013).

- *Natural forcing uncertainty* due to volcanoes, solar cycles, and other natural process arises because we are unable to predict these phenomenon accurately. Throughout this thesis I consider only the climate response to anthropogenic forcing, and so do not consider natural forcings except when calibrating our model to historic observational trends.

In Figure 1.1 Hawkins & Sutton (2009) have plotted the magnitude of the first four sources of uncertainty described above, however they have combined process uncertainty and parameter uncertainty into a single 'model uncertainty'. The figure shows that scenario uncertainty is by far the dominant source of uncertainty in projections of global-mean decadal-mean warming by the end of this century.

1.1.1 Scenario uncertainty

Much research is currently being done in the field of climate prediction, and most of it can be described as quantifying or reducing one or more of these types of uncertainty. The research that I am undertaking as part of my DPhil lies at the interface between scenario and model uncertainty.

Until recently, scenario uncertainty was the field of social scientists and the Intergovernmental Panel on Climate Change (IPCC) Working Group 3 (WG3). This started to change in 2009 with a series of papers introducing the concept and features of cumulative carbon emissions which will be introduced in full in section 1.2.2 (Anderson 2008; Allen *et al.* 2009c; Matthews *et al.* 2009; Meinshausen *et al.* 2009; Zickfeld *et al.* 2009). This concept allows for an innovative method of simplifying scenario uncertainty, as these papers show that when considering cumulative emissions, warming is path-independent. Additionally, Allen *et al.* (2009c) suggest that the uncertainty in warming is also path independent for a given cumulative emissions total, a concept I will expand in Chapter 5 of this thesis. I

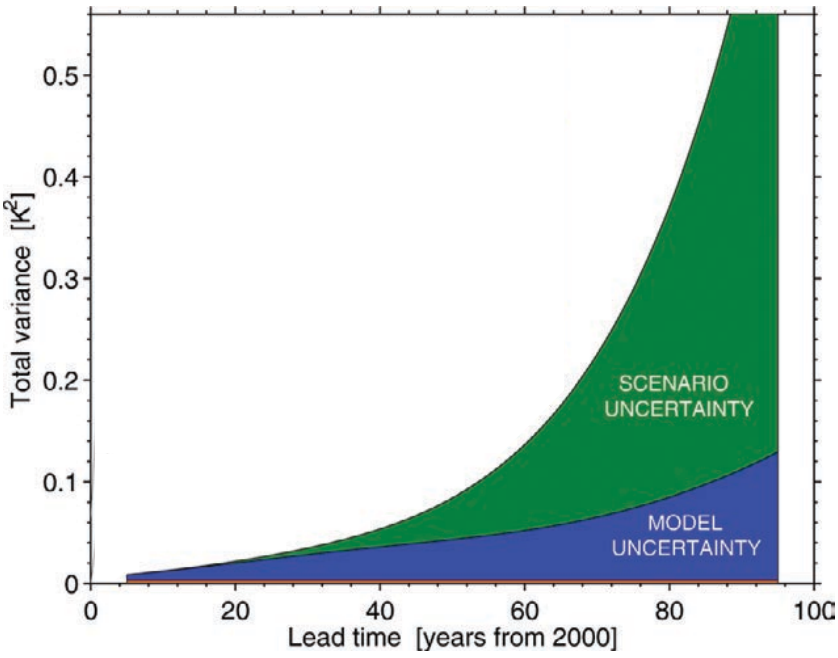


Figure 1.1: Total variance for global-mean, decadal mean surface air temperature predictions taken from 15 climate models and three emissions scenarios, split into the three sources of uncertainty. Green represents scenario uncertainty, blue is model uncertainty, and orange is uncertainty due to internal variability. Initial condition uncertainty is calculated from the internal variability of each model. Model uncertainty is calculated from the spread of model results from the 15 simulations of each emissions scenario. Scenario uncertainty is calculated as the spread between ensemble means for the different emissions scenario. Source: Hawkins & Sutton (2009)

will be using the concept of cumulative carbon emissions throughout this thesis to explore scenario uncertainty from this new angle.

I am currently contributing to work with Le Quéré *et al.* (2013) on quantifying and reducing scenario uncertainty directly, which I do not discuss in this thesis. The work presented in this thesis is concerned with exploring this uncertainty in more detail and determining the emissions targets necessary to avoid dangerous climate change. I will explore different definitions of dangerous throughout this thesis, but the most common will be 2°C of warming above pre-industrial levels, the origins of which I discus in Section 1.3.1.

As the largest contributor to historic global warming, and to the scenario uncertainty in future warming, let us begin with a discussion of carbon dioxide (CO_2).

1.2 Cumulative carbon emissions

1.2.1 CO_2 in the distant future

A substantial fraction of the CO_2 released into the atmosphere by human activity remains there, in effect, for centuries to millennia (Archer 2005). Changes in ocean chemistry, which can be described through the Revelle buffer factor (discussed in Section 2.3.3) limit oceanic removal of CO_2 (Solomon *et al.* 2009) while the potential for terrestrial vegetation to take up CO_2 is also predicted to fall as the climate warms (Cox *et al.* 2000), although the size of this feedback is uncertain (Friedlingstein *et al.* 2006). Complete removal requires geological timescales, (Archer 2005; Archer & Brovkin 2008) or assistance from large-scale air capture technologies (Keith *et al.* 2006; Keller *et al.* 2008).

This feature of the climate implies that bringing future emissions to zero would not reduce temperatures except in the very long term, but would rather hold temperatures almost steady (Matthews & Caldeira 2008; Matthews & Weaver 2010). Several recent studies have sought to exploit this observation in order to provide a simple link between future warming

and the integral of emissions over time, a quantity often referred to by the term *cumulative carbon emissions* (Allen *et al.* 2009c; Matthews *et al.* 2009; Meinshausen *et al.* 2009; Zickfeld *et al.* 2009).

1.2.2 Introducing cumulative carbon emissions

Three studies were published in the journal *Nature* in 2009 that illustrated a link between cumulative emissions of CO_2 and global mean warming (Allen *et al.* 2009c; Matthews *et al.* 2009; Meinshausen *et al.* 2009). Cumulative emissions are the sum of emissions over time, which is the equivalent of the integral under an emissions pathway such as the three illustrated in Figure 1.2. Allen *et al.* (2009c), considered the cumulative carbon emissions summed between pre-industrial times and 2500, linking them to peak warming. Meinshausen *et al.* (2009), examined multi gas pathways and used a cumulative emissions metric between years 2000 and 2050 to relate to the probability of exceeding a 2°C target, rather than the amount of warming. The German Advisory Council on Global Change (WBGU 2009) argued for a cumulative limit between 2010 and 2050, and Matthews *et al.* (2009) argued that warming by a given date is proportional to cumulative emissions to that date, and hence peak warming is determined by cumulative emissions up until the time of peak warming.

These papers show how cumulative emissions provide a tractable, well-constrained and concise metric for use by policy makers interested in avoiding some level of peak global warming (England *et al.* 2009). The recent Copenhagen Accord (UNFCCC 2009) contains an aim of limiting warming to no more than 2°C, and draws on earlier targets from the EU and G8 (den Elzen & Meinshausen 2006; G8 2009). Though not specified in the Copenhagen Accord, this 2°C warming limit is usually presumed to be relative to pre-industrial levels (Rogelj *et al.* 2010), a point that was confirmed in the UN negotiations in Cancun (UNFCCC 2010). Using the results in Allen *et al.* (2009c), a 2°C limit on the most likely peak CO_2-induced warming could be achieved by limiting cumulative emissions to one trillion tonnes of carbon (1 TtC). By comparing the estimates from Allen *et al.* (2009c),

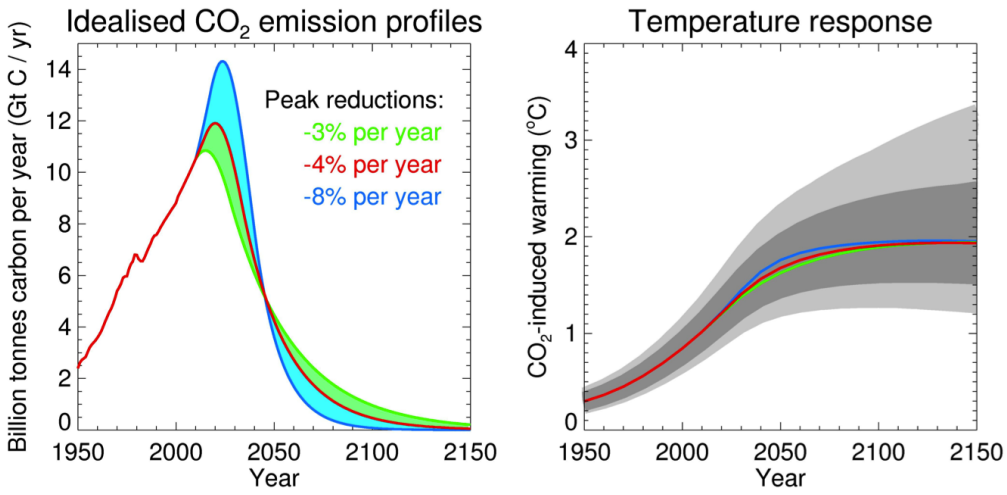


Figure 1.2: Left panel shows idealised CO_2 emissions pathways, each consistent with total cumulative emissions of 1 trillion tonnes of carbon (1 TtC). Right panel shows the temperature response for each of these pathways. Grey shading, representing the likely and very likely range, is explained in section 2.6.2. Source: Allen *et al.* (2009b)

Meinshausen *et al.* (2009), Matthews *et al.* (2009) and Zickfeld *et al.* (2009) we see that there is broad agreement that approximately one trillion tonnes of carbon (1 TtC) causes a best guess warming at the time of peak temperature of 2°C, though the range on each estimate is quite large. More recent estimates of the transient climate response to emission (TCRE) by Johns *et al.* (2011), Rogelj *et al.* (2012a), Eby *et al.* (2013) and Gillett *et al.* (2013) suggest the range is between $0.8°CTtC^{-1}$ and $3.0°CTtC^{-1}$, with a best estimate in the $1.5–2°CTtC^{-1}$ range (Collins *et al.* 2013). These estimates are discussed in more detail in Chapter 8.

Cumulative emission targets represent the sum of emissions over time, and therefore these cumulative emissions could be distributed over time in a number of ways (Allen *et al.* 2009c). For example, an early peak in emissions could be followed by a relatively slow rate of post-peak decline, or a later peak followed by a much more rapid decline (Allen *et al.* 2009b). This is illustrated in Figure 1.2, in which the three emissions pathways have the same cumulative total and therefore resultant warming, despite having different peak dates and rates of decline. One real-world difference between the pathways is that it may not be

technically or politically feasible, or economically desirable, to decrease emissions at rates much in excess of 3 or 4% per year so that peaking later may not be viable (Davis *et al.* 2010; Friedlingstein *et al.* 2011; Meinshausen *et al.* 2009; Mignone *et al.* 2008; Rive *et al.* 2007).

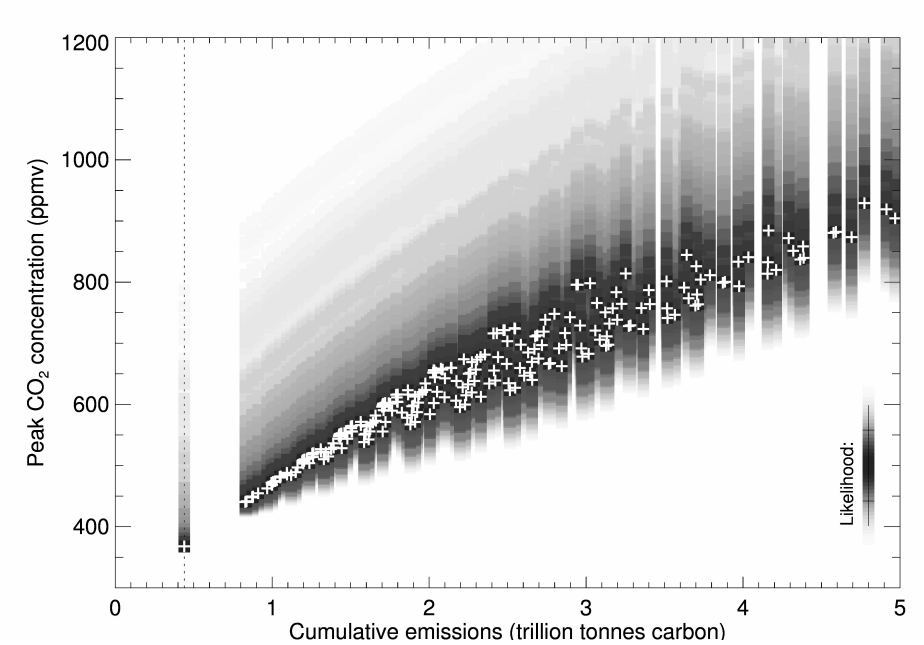
1.2.3 Concentrations, temperatures and cumulative emissions

Having established that cumulative emission should correlate with best guess temperatures, we can now explore this correlation, including parameter uncertainty, and also the comparisons between cumulative emissions and concentrations. This has been done in Figure 1.3. Figure 1.3b is Figure 1 from Allen *et al.* (2009c) with the coloured lines and points removed, and Figure 1.3a is the same plot but for peak CO_2 concentration, rather than peak warming.

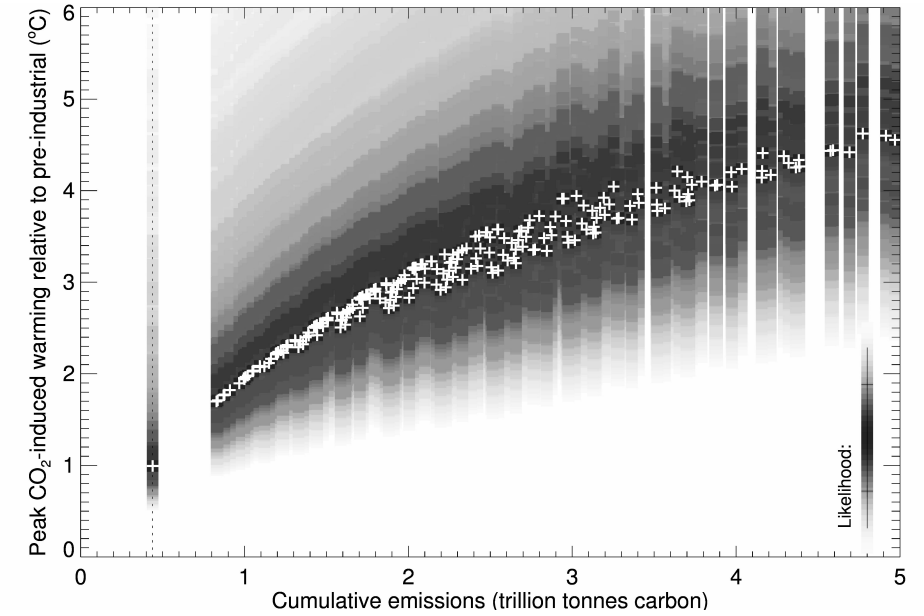
In Figure 1.3 we see how the uncertainty on the temperature and concentration response to a given cumulative total (grey shading), is much more than the spread in best guess values for the range of scenarios (white crosses). The spread in the white crosses is purely due to scenario uncertainty, but the spread in the grey shading is also due to parameter uncertainty.

The spread of the parameter uncertainty in the concentrations plot is less than the spread of parameter uncertainty in the temperature plot. This difference in uncertainty is because the concentrations plot only considers uncertainty arising from the carbon cycle, while the temperature plot considers both uncertainty from the carbon cycle and uncertainty from the climate response.

The spread in best guess values (white crosses) for concentrations in Figure 1.3a is wider than the spread of best guess temperatures in Figure 1.3b. We can conceptualise this difference by considering the differences between the response to an emissions pulse and a sustained emission over many decades with the same cumulative total. As shown in Figure 2.3, the pulse emission will result in carbon concentrations rising instantaneously, then falling rapidly, and then decreasing much slower. The temperature response to the



(a) Peak CO_2 concentrations and cumulative emissions



(b) Peak warming and cumulative emissions

Figure 1.3: (a) peak CO_2 concentration and (b) peak CO_2-induced warming for 250 idealised emission pathways as a function of total cumulative emissions 1750-2500. Figure 1.2 contains three of the emissions pathways, and three others are plotted in black in Figure 5.1. White crosses correspond to best-fit values of simple climate model parameters, with each cross corresponding to a single scenario. Grey shading shows relative likelihood of other parameter combinations, plotted in order of increasing likelihood, showing the uncertainty in peak warming arising from parameter uncertainty in the simple model, as explained in section 2.6.2. Colours used are the same as in those in the plume in Figure 2.11. The uncertainty in (a) includes only carbon cycle uncertainty, while in (b) it includes both carbon cycle and climate uncertainty. The vertical scale bars show the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively. Source: Allen *et al.* (2009c)

pulse emission would increase and then stabilise, at which point the slow ocean equilibration would largely compensate the slow CO_2 draw-down resulting in a roughly constant temperature (Solomon *et al.* 2009). A sustained emission on the other hand would see a near-linear increase of concentrations and temperatures up to a peak value, as we showed in Figure 1 of Smith *et al.* (2012). The pulse emissions would thus have significantly higher peak CO_2 concentration than the sustained emission, but only a slightly higher peak temperature.

One of the reasons why the concept of cumulative emissions has been garnering attention is because it is highly policy relevant. I will now provide an introduction to emissions targets and the history of their use in climate policy before discussing cumulative emissions in the context of climate policy.

1.3 Climate policy

1.3.1 Temperature targets

An aim of the United Nations Framework Convention on Climate Change is "to hold the increase in global average temperature below 2°C or 1.5°C above pre-industrial levels" (UNFCCC 2011). This is an aim that has been agreed by over 100 nations from around the world (UNFCCC 2011). Although current emissions reduction pledges put us on a path with significant risks of exceeding 2°C (UNEP 2012; Höhne *et al.* 2012; Rogelj *et al.* 2011a), the goal is still commonly assessed to be technically possible (UNEP 2012; IPCC 2011; Rogelj *et al.* 2012b).

If society is able to reduce emissions such that they peak soon, so that we can keep temperatures below an overall cumulative total of 0.75–1 TtC, then it is more likely than not that we would be able to keep temperatures below 2°C (Rogelj *et al.* 2012b; Meinshausen *et al.* 2009; Allen *et al.* 2009c). If society's aim was to keep temperatures below, say, 4°C, then we could wait a couple decades before it was necessary to take action to reduce emissions as shown in Figure 7.7c (Rogelj *et al.* 2011b, 2012a). However, to meet a two

degree goal, emissions must peak soon, because of technological, economic and political constraints on the maximum possible rate of post-peak emissions decline (Davis *et al.* 2010; Friedlingstein *et al.* 2011; Meinshausen *et al.* 2009; Mignone *et al.* 2008; Rive *et al.* 2007) This urgency is a motivation for this thesis, and a reason why I am researching emission targets to avoid dangerous climate change.

1.3.2 Concentration targets

The aim of keeping temperatures below 1.5 or 2°C is relatively new, and prior to this the stated aim of the UNFCCC was, and technically still is, the "stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system" (United Nations 1992).

Exactly what qualifies as 'dangerous' anthropogenic climate change is not a question that can be answered objectively without a value judgement (Smith *et al.* 2009; Schneider 2001). The impacts of climate change are both global and regional, and while most impacts will be negative some, such as the boost to agricultural yields at higher latitudes, are likely to be positive, at least for smaller levels of temperature increase (Smith *et al.* 2009; Easterling & Aggarwal 2007).

For the purposes of this thesis, I will use the definition of dangerous a 1.5 or 2°C global mean temperature increase relative to pre-industrial levels. In this thesis I will focus more on 2°C than 1.5°C, as restricting temperature increases to 2°C is the goal of my sponsors the UK Government (UK Government 2012b), and also because it is considered to be more technically feasible than the 1.5°C goal (Davis *et al.* 2010; UNEP 2012).

1.3.3 Emission targets

For a long time policymakers have been considering different types of emissions targets aimed at slowing the rise of global mean temperatures (United Nations 1992). In 1992 the United Nations Framework Convention on Climate Change (UNFCCC) was created in order to facilitate negotiations on these emissions reductions (United Nations 1992). In

1997 the UNFCCC agreed on the Kyoto Protocol, which was brought into effect in 2005, representing an agreement among nations to reduce emissions by percentages named in the text by 2008–2012 relative to 1990 levels (United Nations 1998). This is an example of emissions targets that state that emissions will be at or below a certain level in a certain year or window of years (e.g. 2008–2012). More targets of this type have been negotiated in the UNFCCC since (UNFCCC 2009, 2010, 2011). The benefit of emissions targets that consist of limiting emissions to a certain rate at a certain date is that they are easy to understand, and particularly when the date of the target is set near enough to the present they can encourage governments to take action on reducing emissions (United Nations 1998; UNFCCC 2009). The downside is that they do not provide a strong constraint on the cumulative total carbon emissions, which is ultimately what impacts temperatures. I compare this type of target with cumulative targets quantitatively in Chapter 4.

Another type of target that has been discussed, particularly in the non-governmental organisation (NGO) community, targets the date at which emissions peak (Huntingford *et al.* 2012). This type of target illustrates the scale of the problem at hand, namely shifting the trend in emissions from an exponentially increasing one to a trend with a negative gradient, but provides a weak constrain on cumulative emissions and hence the magnitude of warming as I show in Chapter 4.

Since the high-profile papers on cumulative emissions were published, policy has shifted somewhat more towards the direction of using cumulative targets (WBGU 2009), though it is still emissions at a given date targets that are mostly discussed at the UNFCCC (UNFCCC 2009, 2010, 2011).

1.3.4 Cumulative emissions in climate policy

There have been a number of recent climate policy proposals that aim to limit cumulative emissions. One, by the German Advisory Council on Global Change (WBGU), is to constrain cumulative emissions between now and 2050 (WBGU 2009), following Meinshausen

et al. (2009). In Chapter 4 I compare a cumulative target to 2050 against a cumulative target to a later date in terms of their ability to constrain peak warming.

A similar proposal by Kallbekken *et al.* (2009) aims to constrain cumulative emissions to 2050 in order to constrain peak warming, but also to constrain cumulative emissions to 2030 in order to constrain the near-term rate of warming.

The UK Government has taken a hybrid approach and used short-term budgets over five-year periods (UK Government 2012a). The UK Government has not, to the best of my knowledge, set an overall cumulative budget though, but rather a set of short-term budgets (UK Government 2012a).

Kyoto 2 (Tickell 2008), SAFE Carbon (Allen *et al.* 2009b) and Contraction & Convergence (Stott 2006) are proposals that constrain cumulative targets through a different set of policy mechanisms. Kyoto 2, which is unrelated to the Kyoto Protocol, aims to limit cumulative emissions by capping and trading the amount of carbon, in the form of coal, oil and gas, that can be removed from the ground (Tickell 2008).

SAFE Carbon, instead of limiting the amount of carbon that can come out the ground, uses mandatory sequestration to force an increasing fraction of carbon burned to be sequestered (Allen *et al.* 2009a). The fraction of carbon that has to be sequestered would approach 100% as the total carbon emitted approaches the cumulative limit.

Contraction & Convergence divides the cumulative emissions between nations by setting national cumulative emissions budgets (Stott 2006).

Though cumulative carbon emissions provide a strong constraint on warming, putting an emissions target based on cumulative emissions in to policy would have some potential problems associated with it.

1.3.5 Policy limitations of cumulative emissions

Because cumulative emissions targets are long-term emissions targets, they allow policy-makers to be unconstrained by the target for a long time. For example, despite the fact that society has burned less than two thirds of a 1TtC cumulative budget, a 1TtC global

emissions target would require emissions to peak soon if unrealistic rates of emissions reductions of more than a few percent per year are to be avoided (Rogelj *et al.* 2011b; Allen *et al.* 2009c, b). Instead with a cumulative target one generation of policymakers is able to chose to leave the emissions reductions to a future generation of policymakers with out suffering the consequences of exceeding the target. This 'kicking the can down the road' is a common problem with many types of long-term target (Bosetti & Victor 2011).

A related problem arises from investors not trusting future governments to stick to a target, and thus investing less in industries that would allow emissions to fall (Fuss *et al.* 2009; Bosetti & Victor 2011). Governments can in theory take steps to reduce the scale of this credibility problem, for example by issuing 'carbon bonds' that only pay out if Government is unable to stay within a cumulative emissions budget (Hepburn 2012). These carbon bonds would allow investors to hedge against the risk of a cumulative emissions target not being met and the affect of such a failure on low-carbon industries. These carbon bonds would also serve as a financial incentive for future Governments to meet their cumulative emission targets. Additionally, these bonds would act as a type of prediction market on the chance of a Government meeting their cumulative emissions target, allowing an electorate to see the impact of a Government's policy changes on the chance of meeting that cumulative emissions target in real time. As with all Governments bonds, investors would be running the risk of a Government default, and this risk would be priced into the cost of the bond. These carbon bonds could be applied to any type of emissions target, and are one of the few proposed mechanisms for reducing the impact of the credibility problem in setting long-term emissions targets.

A further problem with a credible cumulative emissions target is that it would convert the right to emit carbon from a non-rivalrous, non-excludable 'public good' into a rivalrous, non-excludable 'common good' (Ostrom 2003). This would encourage nations to extract and sell their carbon reserves as soon as possible before the emissions cap is reached. This problem can be partially avoided by treating the right to emit like an excludable 'private

good' by negotiating cumulative budgets for each nation, nations setting cumulative budgets for sub-national actors, and then optionally allowing trading between actors (Ostrom 2003). While this approach would help with a Kyoto 2 or the Contraction & Convergence framework, it would not help avoid a common goods problem if a SAFE Carbon approach were taken because in this proposal actors are not limited in the total amount of carbon they can extract, only in the fraction they must sequester (Allen *et al.* 2009a; Ostrom 2003).

In order to model the impact of cumulative emissions on the climate, we need to use a climate model. In the next sections, and in the following two chapters, I will motivate, introduce and then test the climate model that I use in this thesis.

1.4 Modelling global warming

1.4.1 Model hierarchy

The world is a complicated place. The phenomena that govern global climate change range from Rossby waves with characteristic length scales of the order of 1000 km (Rossby 1938), to cloud physics that occurs on microscopic scales (Baker 1997). The timescales over which processes relevant to global climate change occur range from microseconds for phenomena like contact freezing in precipitation formation to millennia for carbon uptake due to geological weathering. Climate physicists attempt to model this complexity by representing these phenomena with equations (Stocker 2011)[1]. These equations are usually too complex to be solved analytically, and so must be solved numerically using computer programs that are generally referred to as climate models.

We are unable to represent the full complexity of the climate system in our models as they would be too computationally-intensive to run on present-day computers. Thus climate modellers have to represent the world incompletely in their models. More complex climate models represent the world as a grid, and model the interactions between grid

[1] Stocker (2011) provides a good introduction to climate modelling including a discussion of simple models and their pros and cons, and provides the background on many of the topics I will discuss in the rest of this chapter. To avoid repetition I will not repeatedly reference Stocker (2011) throughout the remainder of this chapter.

boxes. Due to computational restrictions, grid boxes are often of the order of 100 km wide, and so there are many sub-grid scale phenomena that cannot be represented using this methodology. These phenomena are often represented using parameterisations that model them approximately rather than explicitly. For example we are currently unable to represent clouds explicitly in our models due to their small scales, and so instead we make simplified parameterisations of their behaviour that use a number of parameters such as the raindrop descent rate or the entrainment coefficient that determines how much air mixes as it rises. Phenomena that occur on larger scales, such as the large-scale flow of the atmosphere, are represented explicitly in most general circulation models (GCMs) that use the grid box methodology.

The differing levels of complexity in different climate models lead to a hierarchy of models that have different attributes and are useful for different things. On one end, state-of-the-art coupled atmosphere-ocean general circulation models (AOGCMs) contain atmospheric chemistry, fully dynamic oceans, complex ice physics, an interactive land surface, and many other phenomena. These GCMs contain hundreds of thousands of lines of code, tens of thousands of grid boxes, and hundreds of equations to represent the climate system. At the other end of the hierarchy, simple box models use only a handful of equations to represent the climate system.

1.4.2 Physics of mitigation

Three of the most basic physical components of climate change mitigation are the carbon cycle, radiative forcing, and the climate response. Modelling the carbon cycle entails simulating the flow of carbon between different pools, including pools in the ocean, atmosphere, and biosphere, which are discussed in more detail in Chapter 2. The carbon in the atmosphere exerts a radiative forcing, which can be crudely thought of as absorbing and re-emitting more long-wave radiation than would otherwise occur, increasing the effective emission height at which radiation escapes to space (Manabe & Strickler 1964). This is what increases radiative forcing, the net change in incoming minus outgoing radiation at

the tropopause due to greenhouse gases (GHGs) (Planton 2013). This radiative forcing influences the global energy balance, increasing the thermal energy in the Earth system, and slowly increasing global mean temperatures following the timescale of the climate response, as determined by the global effective heat capacity, as discussed in more detail in Chapter 2. These three physical components are modelled crudely in even the most basic coupled climate-carbon cycle simple models.

Most simple models assume that these three components of the climate system are approximately linear, because approximate linearity emerges as a feature of most GCMs (Matthews *et al.* 2009; Gillett *et al.* 2013). There could however be non-linearities in the climate-carbon cycle system, as discussed in Chapters 2 and 3, which have the potential to raise temperatures higher than would be projected in a linear model. There are also many other systems that influence the global mean climate that I have not discussed here, such as the cryosphere and ocean dynamics, which are driven in turn in part via emissions by human systems such as the economy and the political system.

1.4.3 Modelling global mean temperature in simple models

In this thesis, I will almost exclusively model global mean temperatures in a simple coupled climate-carbon cycle model that I will introduce in Chapter 2. I can use a simple model because simple models can be made to mimic GCMs when tuned appropriately, as I discuss in Section 1.5.1.

If I were attempting to model regional warming, or a regional climate impact, I would probably need a significantly more complex model than the simple model that I will use in this thesis. The benefit of modelling global mean temperatures is that we can treat the Earth as a single system and only consider energy fluxes in and out (Budyko 1969). Additionally, CO_2 is a well-mixed GHG, and so I can consider global mean CO_2 concentrations, and total annual carbon emissions.

Many of the sub-systems that make up the climate system are non-linear, such as water vapour feedbacks or ice albedo feedbacks, however the climate system is so large with so

many diverse components that when we consider the impact of all of these sub-systems in aggregate we find that the climate system behaves approximately linearly (Christensen & Kanikicharla 2013; Matthews *et al.* 2009).

Using a simple model means that we will not be able to simulate the dynamic behaviour of the climate system, and in particularly the inter-annual variability displayed by the global mean temperature (Stone 2005; Jones *et al.* 2012; Morice *et al.* 2012). Because we do not include this dynamic behaviour in the simple model we can only model temperature trends, rather than the far noisier actual global mean annual mean temperatures (McAvaney *et al.* 2001). As a result we are unable to simulate phenomenon like the recent warming 'hiatus' (Kosaka & Xie 2013; Guemas *et al.* 2013).

In the next section I address the question of why I focus on global mean temperatures, and in particular peak temperature despite the fact that peak temperature is only one of many indicators and impacts of climate change, as I discuss in Section 3.7.1.

1.4.4 Peak warming

Much of this thesis is devoted to analysing the impact of various climate mitigation measures and metrics on peak warming. A large motivation for this focus on peak warming is due to the fact, discussed in Section 1.2.1, that once temperatures rise they are unlikely to fall again due to the likely irreversibility in the climate system, unless society in engages in large scale geoengineering through either carbon dioxide removal (CDR) or artificially raising the Earth's albedo through solar radiation management (SRM) (Lowe *et al.* 2009; Solomon *et al.* 2009; Vaughan & Lenton 2012). Thus the peak temperature is approximately the temperature that society is likely to have to live with for decades or even centuries after temperatures peak (Lowe *et al.* 2009; Solomon *et al.* 2009).

Additionally, in most climate models, many impacts scale approximately with global mean temperature, a phenomenon known as 'pattern scaling' (Mitchell 2003). Though pattern scaling is not perfect, it does predict many of the major features of future climate impacts that scale with temperature, such as precipitation and regional temperature changes

(Mitchell 2003). Because climate impacts scale approximately with temperature, the total impact may scale as a function of peak warming, particularly when considering the potential irreversibility of this temperature rise (Fankhauser *et al.* 2001).

There are also political reasons for focusing on peak warming, such as the fact that an aim of the UNFCCC is to limit peak warming to 1.5 or 2°C as discussed in Section 1.3.1. Peak warming also serves as a climate change impact for the UNFCCC to negotiate on and ultimately attempt to avoid that has a greater level of predictability than for example avoiding a certain level of Arctic sea ice loss, but is also closer to actual climate change damages in the climate change cause-effect chain shown in Figure 3.7 than say atmospheric CO_2 concentrations (Fankhauser *et al.* 2001). Policy-makers could have chosen to avoid a certain level of warming by a certain date, for example, but as they have instead set peak warming targets. I hope this thesis will contribute to the body of literature that is helping inform the decisions made as a result of these targets.

1.5 Simple models

I will now explain in more detail why I chose to use a simple coupled climate-carbon cycle model throughout my DPhil, despite having far more complex models available to me that could represent the physical processes at play far more realistically.

1.5.1 Benefits of using simple models

I address the limitations of simple models, and in particular the simple model that I am using, in Section 3.7.1, and here outline some of the reasons why I chose to use a simple model in this thesis.

1.5.1.1 Model of a model

Simple models cannot, by definition, represent all of the complexity of more complex GCMs. Instead they can be thought of as a model of a more complex model of the climate-carbon cycle system. Simple climate models are often tuned to the results of more complex

climate models. As a result they are often able to track the trends in the results produced by more complex models. If all we are interested in is trends, then a simple climate model can be as useful as a complex model, provided that we remain within the regime in which the simple model is applicable (see Section 3.7.1 below). The simple model that I am using has a linear temperature response (except for the logarithmic forcing term), an approximately linear carbon cycle response term, but a non-linear temperature feedback on emissions, as discussed in Chapter 2.

1.5.1.2 Insight

Simple models can be extremely useful for developing insight into the fundamental nature of how a system operates. For example, we can easily add and remove components to understand what different parts of the model do. We can also test hypotheses of how different components of the system might work. As an example, we might want to test whether or not the transfer of carbon into the deep ocean can be modelled better as an advective or a diffusive process. For an advective processes, the rate depends on the contrast between the two regions and therefore on the current concentration anomaly, but for a diffusive process the rate depends on the gradient and therefore on the recent history of the anomaly. Thus a simple model could allow us to easily conceptualise how the atmospheric CO_2 concentration response to emissions would change if the transfer of carbon into the deep ocean were purely advective or diffusive.

1.5.1.3 Flexibility

The core of the simple climate model used throughout this thesis contains only a few dozen lines of code. As a result it is possible for a DPhil student to understand the model in its entirety. In contrast, there are probably few people in the world who fully understand the codebase of a GCM in its entirety. Because simple models are so easy to understand they are also easy to edit and use. When we couple this with the speed advantages of using a simple model, we are able to design, code, run, analyse and evaluate an experiment using

a simple model in hours or days, rather than weeks or years. This has allowed me to run perhaps 100 different types of experiment over the course of my DPhil, rather than the much smaller amount that could have been run if I were using a more complex model. Because of the overwhelming number of experiments I was able to run during my DPhil, I will not be presenting all of my results here, only the subset that are most interesting.

1.5.1.4 Computational speed

The hierarchy of climate models spans a huge range of speeds. AOGCMs can take weeks to run a single simulation lasting a couple centuries on even the most powerful supercomputers. At the other end, simple box models such as the one used throughout this report can run thousands of millennial-scale simulations of the global climate in seconds on a desktop computer. In total I have probably run about a billion 750-year simulations of the Earth's climate using thousands of different sets of parameters and thousands of different emissions scenarios. In this thesis I am looking to answer questions such as: what aspects of an emissions scenario are most important for various attributes of a temporal temperature pathway? To do this I need to run perturbed parameter ensembles over ensembles of emission scenarios. I need to run this many simulations in order to better represent the uncertainty in both our parameterisation of the model, and also our uncertainty in future emission scenarios. I would not have been able to run even $\frac{1}{1000}^{\text{th}}$ of the number of simulations if I had been using GCMs, even if I were using distributed computing on a platform such as climateprediction.net (Stainforth *et al.* 2005).

1.5.1.5 Use of constraints

The box-diffusion model uses only five variable parameters, and an additional four fixed parameters, as discussed in Chapter 2. In contrast a GCM might have hundreds of parameters, dozens of which might be variable. As there are not enough observations to constrain all of these parameters in GCMs, expert elicitation is often used to set many of the parameters. Expert elicitation is known to have a number of biases associated with it (Kynn

2008), however this is one of the best method available to GCM designers attempting to select parameters. Simple models on the other hand can be parameterised very differently, provided that there are as many constraints as there are parameters. In the box-diffusion model, we use real-world observations combined with results from state-of-the-art coupled climate-carbon cycle models to determine the values of the variable parameters, rather than the expert elicitation approach. The use of constraints is discussed further in Section 2.6.

1.5.1.6 Parameter uncertainty

In order to quantify uncertainty we need to be able to probe parameter-space to determine how well different sets of parameters reproduce historic temperature trends when used in the model. To properly probe parameter space, the number of runs we need scales with the power of the number of variable parameters. As there are far more parameters in a GCM than in a simple model, this means many more runs are required. Because of the difference in computational speed between GCMs and simple models, the GCM runs also take far longer to run. As a result we are able to probe a far smaller portion of parameter space using a GCM than we are with a simple model. Parameter uncertainty and how we quantify it is discussed in more detail in 2.6.2.

1.5.1.7 Scenario uncertainty

The dominant source of uncertainty in global mean decadal-mean temperature projections in the later half of this century is uncertainty in how much emissions society will produce in the future, as illustrated in Figure 1.1. In this thesis I attempt to cover a range of possible future emissions scenario space, from the most extreme mitigation scenarios to business-as-usual scenarios. These scenarios are explained in more detail in Section 2.5. Because possible future emissions scenario space is so large, I have run thousands of different emissions scenarios, which would be far more difficult using a GCM. By running large ensembles containing hundreds of different emissions pathways, we can begin to analyse trends across emissions pathways. This method allows us to ask questions such

as "what is it about an emissions pathway that controls the resulting peak rate of global mean temperature increase?"

1.5.2 Uses of simple models

Given the various benefits to using a simple model, and the drawbacks discussed in Section 3.7, there are four main ways in which simple climate models have been used historically in the literature, namely: interpreting observations, interpreting GCMs, in scenario and uncertainty analysis, and before GCMs were computationally feasible as the primary means of climate change research.

Simple models are often used to estimate attributes of the climate system from observations. For example Andronova & Schlesinger (2001) uses a simple model to estimate the equilibrium climate sensitivity (ECS) from observations; Forster *et al.* (2006) combines surface temperature observations with radiative forcing estimates and data from the Earth Radiation Budget Experiment (ERBE) in a simple model to estimate the climate feedback parameter, and in turn ECS; and Otto *et al.* (2013b) uses a very simple model relating changes in temperature, heat content, and radiative forcing, in order to estimate the transient climate response (TCR) and ECS. ECS is the temperature increase that would occur after a doubling of CO_2 concentrations relative to pre-industrial once the system is in its long-term equilibrium (Planton 2013). TCR is the smaller temperature increase that occurs at the moment concentrations are doubled relative to pre-industrial in an idealised experiment in which CO_2 concentrations are increased 1% per year (Planton 2013).

As well as interpreting observations, simple models can be used in a similar way to interpret GCMs. For example Sokolov & Stone (1998) used a two-dimensional model to emulate a number of AOGCMs and show that there is considerable uncertainty in the rate of heat penetration into the deep ocean among these AOGCMs. Interpreting complex models using simple models was done more often historically, when computational resources were such that GCM could not even run century-long runs, let alone ensembles. For example,

Schlesinger & Jiang (1990) use a simple model to emulate and then extend 20-year GCM runs in order to estimate the characteristic timescale of the climate system.

When computational resources were even more limited and GCM could barely be run at all, simple models were the main tool used to investigate climate change. For example Budyko's energy balance model (Budyko 1969) was one of the first to suggest that the Earth's climate may have multiple states of equilibrium, alongside similar models produced around the same time by Erikson (1968), and Sellers (1969).

Today the main use of simple models is still as a computationally cheaper way of mimicking a GCM. Simple models are used to explore scenario and model uncertainty in cases where it would be too expensive to do this using GCMs. The Model for the Assessment of Greenhouse gas-Induced Climate Change (MAGICC) model is probably the most popular simple model used today, and has been used in dozens of studies of this type (Wigley 2006; Meinshausen *et al.* 2011a, b). As there are so many studies exploring scenario and model uncertainty using simple models discussed throughout this thesis (Allen & Frame 2007; Allen *et al.* 2009c; Anderson 2008; Anderson *et al.* 2008; Anderson & Bows 2011; Meinshausen *et al.* 2009; Rogelj *et al.* 2011a, 2012a, b; Shine *et al.* 2005, 2007; Smith *et al.* 2012), I will refrain from discussing this issue at length here and will instead give an outline of the topics discussed in this thesis, all of which will be investigated using the simple model I introduce in the next chapter.

1.6 Thesis outline

The simple link between cumulative emissions of carbon and future warming is causing policymakers to rethink the way that emission targets are set. My thesis aims to refine and extend the scientific underpinnings of this new cumulative approach to setting emission targets for avoiding dangerous climate change.

I start by introducing the simple box-diffusion model that I use throughout this thesis in Chapter 2. In Chapter 3, I test the box-diffusion model, illustrating its strengths and

weaknesses, and show that it is suitable for the uses in this thesis.

In Chapter 4 I show that the spread of possible temperatures after meeting a conventional carbon emissions target in a given year is much wider than the spread of resultant temperatures from meeting a cumulative carbon budget. I also discuss the shape of the cumulative emissions curve introduced in white crosses in Figure 1.3b.

Thus far we have considered only peak warming as an indicator of dangerous climate change, however another key indicator is the rate of global warming (Root *et al.* 2003). Kallbekken *et al.* (2009) suggested that under emissions scenarios consistent with avoiding more than 2°C of global warming, the peak rate of CO_2-induced warming is controlled by the cumulative emissions between 2010 and 2030. In the remainder of Chapter 4 I extend this result to make it more broadly applicable by showing that the maximum possible peak rate of CO_2-induced warming is controlled by the peak rate of CO_2 emissions under a wide range of plausible emissions scenarios.

In Chapter 5 I illustrate that the possibility of having a future floor in emissions is not incompatible with a cumulative budget approach, but imposes constraints on the integration period for computing a useful cumulative total. I use emission floors to show that it is cumulative emissions to approximately 2200 that control peak warming in the simple model I use, not cumulative emissions to 2500 as suggested by Allen *et al.* (2009c) or 2050 as used by Meinshausen *et al.* (2009).

The examples above illustrate how the cumulative budgets approach works well for the CO_2-only case. Several recent papers, however, have highlighted problems with incorporating non-CO_2 GHGs into this framework (Cox & Jeffery 2010; Arora *et al.* 2011). During my DPhil I helped to develop a method of incorporating non-CO_2 forcings into the cumulative emissions framework alongside Smith *et al.* (2012). The results from this work are presented in brief in Chapter 7. As part of this work I had to expand the box-diffusion model outlined in Chapter 2 to include non-CO_2. This non-CO_2 expansion is explained in Chapter 6.

Recently there has been a series of papers published proposing to help tackle climate change by cutting short-lived climate forcers (SLCFs) (Molina *et al.* 2009; UNEP & WMO 2011; Shindell *et al.* 2012). In Chapter 7 I illustrate why delaying the implementation of the measures suggested in these papers until after carbon emissions have peaked, rather than implementing them now, will have little impact on peak warming. I also show that the proposed measures cannot be used as a way of "buying time" to cut emissions of CO_2 if the aim is to reduce peak global warming.

In Chapter 8 I introduce a climate policy based on the work in this thesis and use it to illustrate the benefits and limitations of my work.

In summary, over the course of my DPhil, I hope to refine the scientific underpinnings of the link between cumulative carbon emissions targets and peak global warming, as well as extending the concept to include non-CO_2 GHGs and SLCFs. Over the course of this work, I hope to improve our understanding of the emission targets required to avoid dangerous climate change.

1.7 Relation to published work

This thesis is comprised mostly of papers and a book chapter which are either published or have been submitted for publication. Chapters 1, 2, 4 and 5 have been summarised in a book chapter titled "The importance of limiting cumulative emissions" that was commissioned by *Springer* as part of a book titled "Can we still avoid dangerous climate change?", which is currently in press. This book chapter was written by myself, David Frame, Chris Huntingford, Jason Lowe, Laila Gohar and Myles Allen (Bowerman *et al.* 2013a). Chapters 2, 4 and 5 appeared as an article in *Philosophical Transactions of the Royal Society A*, which was written by myself, David Frame, Chris Huntingford, Jason Lowe and Myles Allen (Bowerman *et al.* 2011). A paper based on the work in Chapter 7 has been accepted for publication by *Nature Climate Change*, and was written by myself, David Frame, Chris Huntingford, Jason Lowe, Stephen Smith, and Myles Allen (Bowerman *et al.* 2013b). I was

lead author on each of these works, and I wrote the text, although I gratefully acknowledge guidance, discussion, comments, and suggested amendments from my coauthors and several anonymous reviewers. I was responsible for all of the modelling, though the CO_2-only box-diffusion model as presented in Chapter 2 was developed by Myles Allen and David Frame.

I am also a co-author on a few papers closely related to the work presented in this thesis. The results distinguishing between SLCFs and long-lived climate forcers (LLCFs) presented at the start of Chapter 7 were published as part of a larger paper in *Nature Climate Change*, written by Stephen Smith, Jason Lowe, myself, Laila Gohar, Chris Huntingford and Myles Allen (Smith *et al.* 2012). My contribution to this paper is discussed in Chapter 7. A follow-on paper to Chapters 4 and 5 using a similar methodology but a different model was published in *Environmental Research Letters*, and was written by Chris Huntingford, Jason Lowe, Laila Gohar, myself, Myles Allen, Sarah Raper, and Stephen Smith (Huntingford *et al.* 2012). This paper is discussed briefly in Chapters 4 and 5. A follow-on paper to Chapters 6 and 7 on metrics, SLCFs and LLCFs is under review at *Environmental Research Letters* and was written by Chris Huntingford, Jason Lowe, Nicholas Howarth, myself, Laila Gohar, Alexander Lorenz, David Lee, Stephen Smith, Michel den Elzen, Detlef van Vuuren and Myles Allen (Huntingford *et al.* 2013). The model used in Huntingford *et al.* (2013) is the multi-gas extension to the box-diffusion model that I developed in Chapter 6. I am also a co-author on Otto *et al.* (2013a) and Le Quéré *et al.* (2013), two related papers which I worked on during this thesis.

Chapter 2

Modelling the coupled carbon cycle-climate system

In this chapter I introduce the simple model I will use throughout this thesis. In Chapter 1 I discussed cumulative emissions and their impact on climate policy, and also discussed simple climate and carbon cycle models. I concluded that in this thesis I would use a simple climate model to explore aspects of the emissions targets required to avoid dangerous climate change, and in particular the role of cumulative emissions and short-lived climate forcers.

This chapter begins by outlining the simple 'box diffusion' coupled carbon cycle-climate model that I will use through this thesis. I will illustrate how the box-diffusion model is related to another very simple climate and carbon cycle models. In the later part of this chapter I will discuss the emissions pathways that I use in this thesis, before moving into a discussion of uncertainty and in particularly the maximum likelihood approach.

First I will introduce the carbon cycle, and consider how to model it simply.

2.1 Carbon cycle

The extent to which the atmospheric composition is altered by the combustion of fossil fuels and carbon from other sources is determined by the global carbon cycle. Figure 2.1 illustrates the flows of carbon in the global carbon cycle during the 1990s, and how these fluxes have changed since the pre-industrial era. The majority of the extra carbon in the

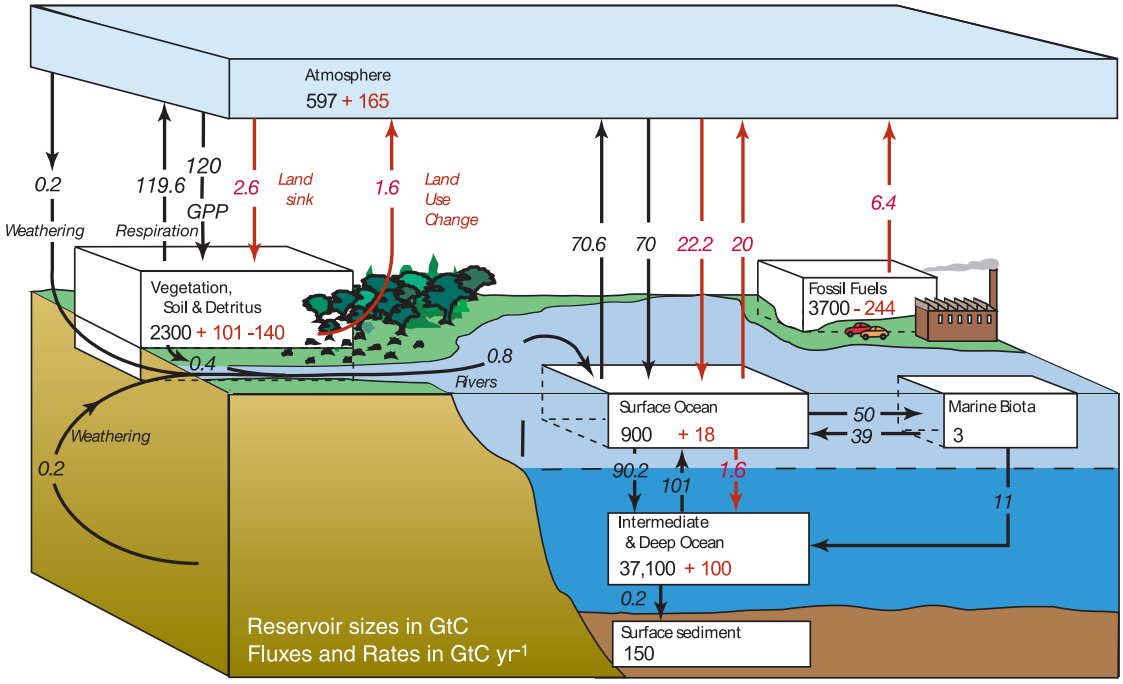


Figure 2.1: The global carbon cycle. Figures shown are the main annual fluxes in GtC yr^{-1} as estimated in the 1990s. Pre-industrial natural fluxes are in black and anthropogenic fluxes are in red. Source: IPCC (Denman *et al.* 2007)

atmosphere has come from the burning of fossil fuels, and a significant fraction also comes from changes in land use. Much of this carbon is ultimately stored in the atmosphere, biosphere, and ocean.

The global carbon cycle is far more complicated than is made out in Figure 2.1, but since we are concerned only with the global mean annual mean atmospheric CO_2 concentration because CO_2 is a relatively well-mixed GHG (Stocker 2011), we can average over much of this complexity and tune a simple carbon-cycle model to the results of the state-of-the-art coupled climate-carbon cycle models that do represent some of this complexity, as explained in Section 2.6.2.

For example, atmospheric CO_2 concentrations have a seasonal cycle because there is more land and vegetation in the Northern hemisphere than in the Southern hemisphere (Etheridge *et al.* 1998a). However, the box-diffusion model has a one-year time step and so we are considered only with annual mean concentrations, and can ignore the seasonal

variation in atmospheric CO_2 concentration.

I will now introduce two simple carbon cycle models that use this approach to simulate the behaviour of the carbon cycle.

2.2 The Bern simple carbon cycle model

In the next section I explain in more detail how the carbon cycle used in the simple model I make use of throughout this thesis works, but first I will introduce a more idealised model of the carbon cycle that has been used throughout the climate change literature (Siegenthaler & Joos 1992; Joos *et al.* 1996, 2001; Godal & Fuglestvedt 2002; Shine *et al.* 2005; Johansson *et al.* 2006; Shine *et al.* 2007; Boucher & Reddy 2008; Boucher *et al.* 2009; Levasseur *et al.* 2010; Fuglestvedt *et al.* 2010; Peters *et al.* 2011b; Johansson 2011; Joos *et al.* 2013). It is a four-box impulse response model, which I will refer to as the Bern simple carbon cycle model. Each of the four boxes in the model takes the form

$$\frac{dC_i}{dt} = \alpha_i E - \frac{C}{\beta_i}, \tag{2.1}$$

where C is the carbon concentration with units of part per million by volume (ppmv), and E is the emissions with units of part per million by volume equivalent (ppmv-eq), the mass of carbon that would have to be injected into the atmosphere in order to instantaneously raise CO_2 concentrations by one ppmv. α_i and β_i are constants as defined in Table 2.1, and i ranges from 0 to 3.

The conversion between ppmv-eq and gigatonnes of carbon (GtC) can be estimated by considering the mass of the atmosphere and the molar mass of air and CO_2. 1 GtC $\approx$ 0.47 ppmv, and thus 1 ppmv-eq $\approx$ 2.12 GtC.

These constants have been obtained by fitting a four-box model of the form of Equation 2.1 with $\beta_0^{-1} = 0$ to the output of an intermediate complexity model with a gridded terrestrial carbon cycle model.

This model can be solved analytically for a wide range of cases, including a pulse injection. For a pulse injection of A ppmv-eq, the atmospheric CO_2 concentration anomaly

i	α_i (unitless)	β_i (years)
0	0.217	∞
1	0.259	172.9
2	0.338	18.51
3	0.186	1.186

Table 2.1: Values for the constants α and β in Equations 2.1 and 2.2. These constants are used in the Bern simple carbon cycle model (Siegenthaler & Joos 1992; Joos *et al.* 1996), as updated by Fuglestvedt *et al.* (2010).

is given by

$$C = A \sum_{i=0}^{3} \alpha_i \exp\left(\frac{-t}{\beta_i}\right), \tag{2.2}$$

where $t \geq 0$ is the time in years after the emissions pulse. This curve is plotted in Figure 2.3 as the red dashed line in each panel.

2.3 A simple box-diffusion carbon cycle model

The model that I have made extensive use of throughout my thesis attempts to mimic the behaviour of the carbon cycle in GCMs using fewer adjustable parameters than are required in an impulse response model. The model was first introduced by Allen *et al.* (2009c), and I explain it in more detail in Bowerman *et al.* (2011). I will introduce each of the equations briefly here before going into more detail about each equation in the following sections.

2.3.1 Carbon cycle model equations

The model uses a 'three box' atmosphere-ocean carbon cycle, in which we assume that the atmospheric CO_2 can be split into three components, the first of which is the long-term equilibrium level of CO_2, C_3. We assume that

$$\frac{dC_3}{dt} = b_3 E \tag{2.3}$$

so that the rate of change of the long-term equilibrium CO_2 level increases with E. $b_3 = 0.15$ is a fixed constant representing the fraction of CO_2 that remains in the atmosphere in the long-term equilibrium (Solomon *et al.* 2009). This fraction is determined in part by the

Revelle Buffer Factor (Archer 2005), which is explained in more detail in Section 2.3.3. C has units of ppmv, and E has units of ppmv equivalent per year.

The second component of the carbon cycle is the exponential advective term, C_2, which is the amount that the CO_2 levels would exceed C_3 if they were ever equilibrated between the atmosphere, land-biosphere and upper-ocean (Allen *et al.* 2009c). C_2 relaxes towards zero through advection into the deep ocean, and takes the form

$$\frac{dC_2}{dt} = b_1 E - b_0 C_2 \tag{2.4}$$

where b_1 is a fixed constant such that $b_1 + b_3 = 0.35$ (Allen *et al.* 2009c). $b_1 + b_3$ represents the fraction of atmospheric CO_2 that would remain in the atmosphere following an injection of carbon in the absence of advection to the deep ocean. b_0 represents an adjustable rate-constant, the inverse of which is of order 200 years (Allen *et al.* 2009c). The use of adjustable parameters is explained in Section 2.6.2. This equation and the next are discussed in more detail in Section 2.3.4.

The third 'box' in our model accounts for the diffusion of CO_2 out of the near-surface mixed layer. The equation takes a similar form to Equation 2.4, except that the final term is replaced by a diffusion term. C_1 is the excess atmospheric CO_2 over the short-term equilibrium level $C_2 + C_3$,

$$\frac{dC_1}{dt} = b_4 E - b_2 \int_0^t \frac{dC_1(t')}{dt'} \frac{dt'}{\sqrt{t - t'}}. \tag{2.5}$$

b_2 represents an adjustable diffusivity, while $b_1 + b_3 + b_4 = 0.85$ is the fraction of CO_2 that would remain in the atmosphere within a year of a pulse injection Allen *et al.* (2009c). $b_1 + b_3 + b_4 < 1$ because some of the CO_2 is rapidly mixed into the ocean mixed layer and biosphere (Hua & Barbetti 2004).

The total concentration of carbon dioxide in the atmosphere, C, is given by

$$C = C_1 + C_2 + C_3. \tag{2.6}$$

The final equation in the carbon cycle model is the temperature feedback term. E, the annual rate of emission of carbon at time t, is given by

$$E(t) = \begin{cases} E_a(t) + b_5 \left(T(t) - \int_0^t T(t') e^{\frac{t'-t}{100}} dt' \right) & \text{for } T(t) > \int_0^t T(t') e^{\frac{t'-t}{100}} dt' \\ E_a(t) & \text{for } T(t) \leq \int_0^t T(t') e^{\frac{t'-t}{100}} dt' \end{cases} \quad (2.7)$$

where E_a is the rate of anthropogenic carbon emission, b_5 is the carbon cycle feedback parameter which is tuned to the results of Coupled Carbon Cycle Climate Model Intercomparison Project (C^4MIP) (Friedlingstein *et al.* 2006) as outlined in Section 2.6.2. $\int_0^t T(t') e^{\frac{t'-t}{100}} dt'$ is simply the exponentially-weighted mean historical temperature, with an e-folding time of one century.

One way of conceptualising this equation is that for smooth emissions curves, when global mean temperatures are increasing the carbon cycle feedback emits additional CO_2 representing emissions from the biosphere. When global mean temperatures are falling we assume that there is no additional emission from the biosphere. In Equation 2.7, there is a period after global temperatures peak when global temperatures are falling, but are still greater than the exponentially-weighted historical mean and so the biosphere is still emitting additional carbon.

For now I will refer to this formulation of the carbon cycle feedback term as the *asymmetric* carbon cycle formulation, as opposed to the *symmetric* carbon cycle formulation which I will introduce in Section 2.3.5.

2.3.2 Conceptualising the simple carbon cycle model

The simple carbon cycle model can be thought of as a simplified version of Figure 2.1, as illustrated in Figure 2.2. Carbon emissions are released into the atmosphere, which has concentration C_α where

$$C_\alpha = C = C_1 + C_2 + C_3. \quad (2.8)$$

Because this model has a one-year time-step, we assume the atmosphere is in equilibrium with the mixed-layer of the ocean and the rapid-response biospheric stores. When

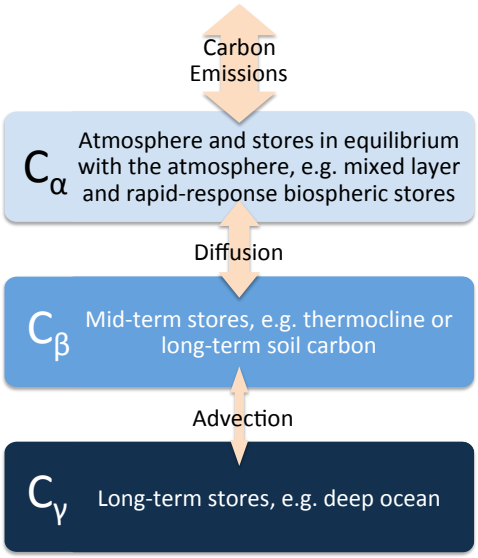


Figure 2.2: A simple schematic of the simple box-diffusion carbon cycle model.

carbon is emitted into the atmosphere, some of this carbon then diffuses into mid-term stores such as the thermocline and the long-term soil carbon. The concentration of CO_2 in these stores is represented by C_β where

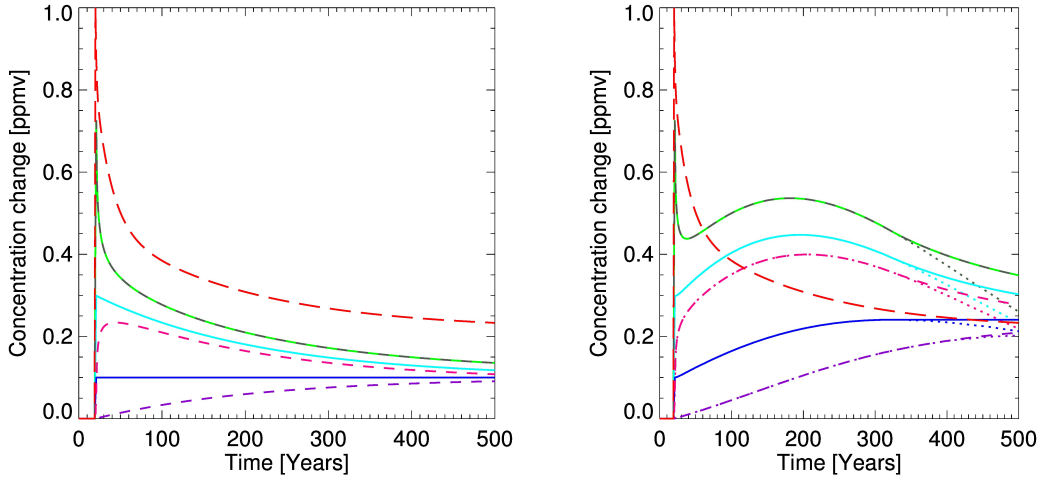
$$C_\beta = C_2 + C_3 - \frac{b_1 + b_3}{b_4} C_1. \tag{2.9}$$

Note that the concept of the 'concentration' of carbon in the thermocline or deep ocean is not as straight forward as it might at first seem due to buffering, as explained in Section 2.3.3. We use it here to mean the concentration of carbon that the air would have in equilibrium if it were in direct contact with the thermocline.

Some of the carbon from the thermocline is then advected into the deep ocean, which has a concentration C_γ where

$$C_\gamma = C_3 - \frac{b_3}{b_1} C_2. \tag{2.10}$$

In Figure 2.3a the concentration of the different boxes and components of the model are plotted before and after a pulse emission of CO_2. We see the concentration of the atmosphere, C_α, rises rapidly and then falls as some of the CO_2 diffuses into the thermocline and mid-to-long-term biospheric store, C_β. In turn, C_β rises as it receives CO_2 from the



(a) Carbon cycle feedback term turned off in the box-diffusion model

(b) Carbon cycle feedback term turned on in the box-diffusion model

Figure 2.3: CO_2 concentration change in different components of the carbon cycle model following a pulse injection of enough CO_2 to instantaneously increase atmospheric concentrations by one ppmv, which requires 2.1 GtC of CO_2. The red long-dashed line is the response from the Bern carbon cycle model as introduced in Section 2.2, using the parameter values from Fuglestvedt *et al.* (2010). The other lines are output from the box-diffusion model run with best guess parameters, as defined in section 2.6.2. C_3 is in dark blue; $C_2 + C_3$ is in light blue; $C_\alpha = C = C_1 + C_2 + C_3$ is in green and grey (upper line); C_β is in pink; C_γ is in purple. C_1, C_2, C_3, C_α, C_β and C_γ are all carbon concentrations, and are explained in Sections 2.3.1 and 2.3.2. The carbon in C_1 diffuses into C_β representing the diffusion of carbon into the land-biosphere, the thermocline, and mid-term carbon stores. The carbon in C_2 advects into C_γ representing the advection of carbon into the deep ocean. Panel a has the carbon cycle feedback term b_5 turned off, i.e. set to zero. In Panel b the carbon cycle feedback is turned on. The solid and dashed lines are simulated using the asymmetric carbon cycle feedback formulation introduced in Equation 2.7, while the dotted lines are simulated using the symmetric carbon cycle formulation introduced in Section 2.3.5. Uncertainty is not included in this figure for clarity, but the uncertainty in the carbon cycle response to a pulse injection of carbon emissions is shown in Figures 3.5, 3.6, 3.8 and 3.9.

atmosphere via diffusion, but then falls as some of its carbon advects into the deep ocean, C_γ. The CO_2 concentration in the deep ocean, C_γ rises slowly and equilibrates at C_3, which is determined by the Revelle buffer factor as described in the next section.

In the box-diffusion model the atmosphere does not increase by one ppmv of CO_2 when I inject enough CO_2 to increase it by one ppmv instantaneously. This is because the model has a one-year time step. We are effectively assuming that about 15% of the CO_2 does not remain in the atmosphere one year after it was injected, consistent with more complex models (Joos *et al.* 2013). Note that the carbon concentration does not increase all the way

up to 0.85 ppmv after the initial injection in Figure 2.3a. This discrepancy is caused by the fact that our pulse is actually a Gaussian pulse over about three years, because a pulse in a single year could cause numerical instabilities in our model. Making this Gaussian pulse wider causes the height of this peak to decrease, and vice versa.

The selection of the parameters used in the model are described in section 2.6. The parameters in this model can be tuned to reproduce the response of the Bern carbon cycle model described in (Forster *et al.* 2007) and based on (Joos *et al.* 2001) by turning the carbon cycle feedback term off.

The behaviour of the model to a pulse of emissions is discussed in more detail in Section 3.5, where I compare the behaviour of the box-diffusion model with the Bern simple carbon cycle model, and a number of earth system models of intermediate complexity (EMICs) and earth system models (ESMs).

2.3.3 Revelle buffer factor

When carbon dioxide is dissolved in the ocean it can partition into one of the components of carbonic acid: dissolved carbon dioxide CO_2^*, carbonate ions CO_3^{2-} or bicarbonate ions HCO_3^- (Williams & Follows 2011). The reversible chemical reactions that govern this process are

$$CO_2^* + H_2O \rightleftharpoons HCO_3^- + H^+, \tag{2.11}$$

$$HCO_3^- \rightleftharpoons CO_3^{2-} + H^+. \tag{2.12}$$

The proportion of the carbon dioxide that partitions into each component depends on, among other things, the pH of the ocean, as illustrated in Figure 2.4. The pH of the ocean is defined as $\text{pH} = -\log [H^+]$, where $[H^+]$ is the concentration of H^+ ions (Williams & Follows 2011). I use the notation [X] to denote the concentration of chemical X. Lower pH indicates a higher $[H^+]$, pushing reactions 2.11 and 2.12 to the left, consuming CO_3^{2-} and favouring CO_2^*.

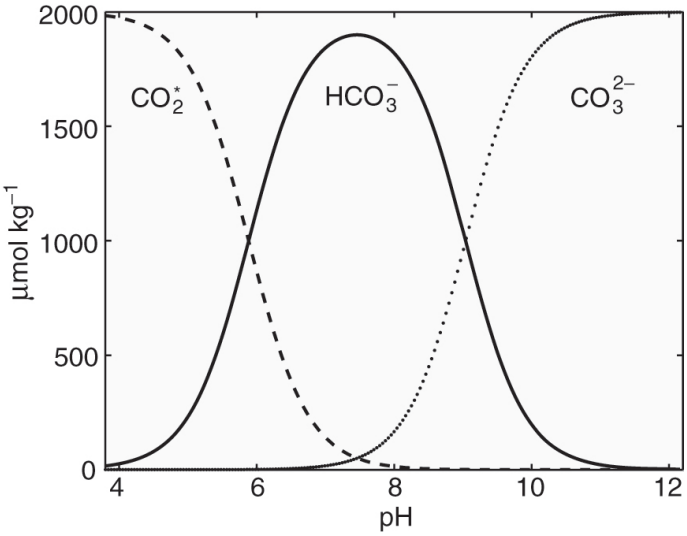


Figure 2.4: The partitioning of carbonate species as a function of pH. Dissolved inorganic carbon, DIC, is partitioned between CO_2^* (dashed line), HCO_3^- (solid line) and CO_3^{2-} (dotted line). Source: Williams & Follows (2011)

As the abundance of H^+ ions is increased or decreased through reactions 2.11 and 2.12, the partitioning of the components of carbonic acid influences the pH of the oceans. Adding more CO_2^* also decreases the pH by allowing more H^+ ions to be created.

The total concentration of carbon in the ocean, or dissolved inorganic carbon, $[\mathrm{DIC}]$, is the sum of the different inorganic carbon species in the ocean (Williams & Follows 2011),

$$[\mathrm{DIC}] = [\mathrm{CO}_2^*] + [\mathrm{HCO}_3^-] + [\mathrm{CO}_3^{2-}]. \tag{2.13}$$

We can describe $[CO_2^*]$ as having partial pressure pCO_2, that is the partial pressure that CO_2 would have in the atmosphere if it were in equilibrium with the ocean while $[CO_2^*]$ is held constant. Now the air-sea exchange of CO_2 depends on pCO_2, and the amount of carbon stored in the ocean is determined by DIC, so let us define a relation between them. We can define, $\bar{B}$, the global mean Revelle buffer factor as

$$\bar{B} = \frac{\frac{\delta \mathrm{pCO}_2}{\mathrm{pCO}_2}}{\frac{\delta[\mathrm{DIC}]}{[\mathrm{DIC}]}} = \frac{\frac{\delta[\mathrm{CO}_2^*]}{[\mathrm{CO}_2^*]}}{\frac{\delta[\mathrm{DIC}]}{[\mathrm{DIC}]}}. \tag{2.14}$$

So $\bar{B}$ is defined as the ratio between the fractional change in $[CO_2^*]$ and the corresponding fractional change in [DIC] (Williams & Follows 2011). Because of the ocean chemistry

described above, changes in DIC are buffered with respect to pCO_2. Or put another way, we can think of there being negative feedbacks in the oceanic chemistry which dampen the impact that changes in pCO_2 have on [DIC].

One way of thinking of the Revelle buffer factor is to consider the counter-example of an inert gas. An inert gas such as CFC-11 would not interact with carbonate chemistry and so in equilibrium a fractional increase in the partial pressure of the gas, pCFC-11, would result in the same fractional increase in the amount of CFC-11 stored in the oceans. Thus CFC-11 would have a buffer factor of unity. One the other hand for CO_2, the change in DIC is about ten times smaller than the fractional change in pCO_2, as defined by $\bar{B}$. Phrased another way, when you add CO_2 to the ocean, the concentration rises $\bar{B}$ times faster than you would expect it would for an inert gas.

However B is not a constant. It varies as a function of, among other things, temperature, salinity and pH. As a result B increases with pCO_2, and as a result DIC is not proportional to pCO_2. B is also higher at higher latitudes. The relationships between these quantities and a map of the spatial distribution of B are given in Figure 2.5.

In this model, in Equation 2.3 we use $b_3 = 0.15$, which was obtained from a fit to more complex models. We can estimate the value of this parameter by considering the fraction of injected carbon that ends up in the ocean in the long term.

$$b_3 = \frac{M_{C_{\mathrm{atm}}}}{M_{C_{\mathrm{ocean}}}} \bar{B}, \tag{2.15}$$

where $M_{C_{\mathrm{atm}}}$ is the mass of carbon in the atmosphere, $M_{C_{\mathrm{ocean}}}$ is the mass of carbon in the ocean. From Figure 2.1, $M_{C_{\mathrm{atm}}} = 760$ GtC, $M_{C_{\mathrm{ocean}}} = 38,000$ GtC. So in the absence of a Revelle buffer factor (i.e. with $\bar{B} = 1$) a GtC pulse would be split between the ocean in the ratio $\frac{760}{38{,}000} = 0.02$. However, as $\bar{B}$ =~9 as shown in Figure 2.5, this estimate gives $0.02 \times 9 = 0.18 \approx 0.2$. 0.2, or 20%, is the value for the long-term fraction of carbon that remains in the atmosphere following a pulse emission obtained by Solomon *et al.* (2009) using AOGCMs and EMICs. It is also approximately the value obtained in the

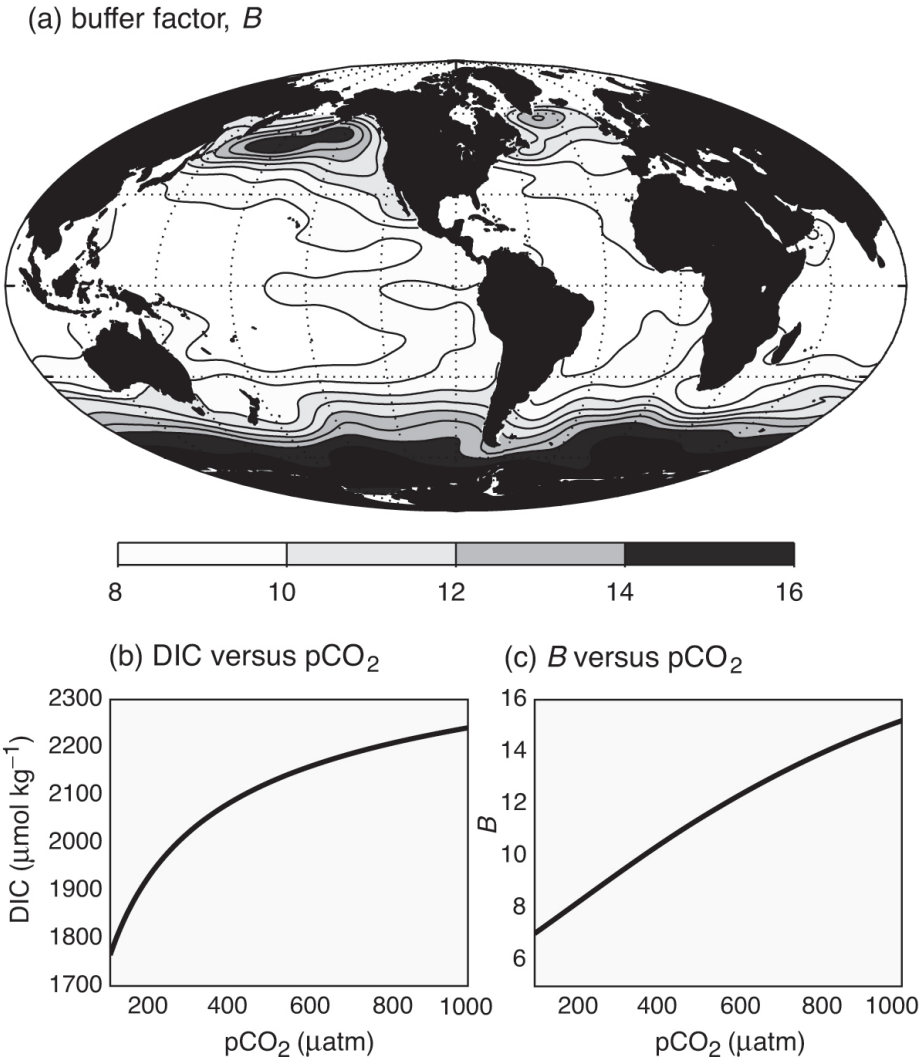


Figure 2.5: Changes in the Revelle buffer factor, B, as a function of (a) location and (c) pCO_2. Panel (b) shows DIC as a function of pCO_2. Source: Williams & Follows (2011)

box-diffusion model for the long-term fraction of carbon remaining in the atmosphere once carbon cycle feedbacks are included, as illustrated by the blue dotted line in Figure 2.3b.

In the model we assume b_3 is a constant, so we assume that changes in b_3 are small compared with the other factors considered here. b_3 has been tuned to reproduce the long-term behaviour of a collection of earth system models of intermediate complexity (EMICs) (Allen *et al.* 2009c; Friedlingstein *et al.* 2006).

2.3.3.1 Limits to this approach

In the box-diffusion model we assume that the Revelle buffer factor, B is a global mean constant $\bar{B}$. In reality, B varies both spatially and temporally (Williams & Follows 2011), and I will now explain how these simplifications impact on our results and how the model could be adapted to represent them more accurately.

Considered purely as a globally averaged quantity $\bar{B}$, the Revelle buffer factor is not a constant as we assume in the box-diffusion model, but instead depends on a number of factors including the global mean temperature T, the global mean salinity, and the global mean alkalinity (Williams & Follows 2011). As the global mean alkalinity is a function of the partial pressure of CO_2 and hence C the atmospheric CO_2 concentration, $\bar{B}$ depends on C. This is illustrated in Figure 2.5. When aggregated, I would expect $\bar{B}$ to increase with increasing C and hence with increasing T. In the box-diffusion model, $\bar{B}$ is a constant and any changes in B at higher C or T will have been incorporated into the the model when it was tuned to more complex EMICs and ESMs,

The Revelle buffer factor, B, varies globally as illustrated in Figure 2.5. The box-diffusion model would need to be adapted to include one or more spatial dimensions to include this inhomogeneity. It is unclear how this simplification impacts on the results presented in this thesis, and the details would probably depend on regional effects and ocean circulations. I would expect modelling B spatially to have a relatively small effect on atmospheric carbon concentrations compared with modelling B as a globally-averaged

dynamic quantity $\bar{B}$. The temporal variations in the spatial distribution would presumably act to increase the temporal variability in the atmosphere carbon concentration C.

Because geochemical processes are not included in this model, we would expect that carbon concentrations in the real world would return to equilibrium faster than they would if purely governed by the Revelle buffer factor. In the box-diffusion model however, the emissions-temperature feedback term returns concentrations back to pre-industrial temperatures unrealistically fast, and so including geochemical processes without first improving the emissions-temperature feedback term would not result in improved performance of the box-diffusion model relative to current EMICs and ESMs.

In the Bern simple carbon cycle model, the lack of a temperature-drive carbon cycle feedback results in a larger Revelle buffer factor once the model is tuned to the output of more complex models.

2.3.4 Advective and diffusive terms

The carbon cycle in the box-diffusion model uses an advective term (Equation 2.4) in conjunction with a diffusive term (Equation 2.5). This was chosen for a number of reasons after discussion with a number of experts. Much of this subsection is the result of discussions between Myles Allen and Chris Jones at the Met Office.

The carbon cycle has elements that are advective, elements that are diffusive, and elements that are neither advective nor diffusive. Because we have five adjustable parameters in the model (constrained by five constraints), and because we are concerned only with global annual averages, we can effectively bundle all of the advective and diffusive terms into just two terms and tune the relevant parameters so that they fit the results of the state of the art coupled climate-carbon-cycle models in C^4MIP. The adjustable parameters and how values were selected is discussed in more detail in Section 2.6.2, and this simplification of the system is discussed in more detail in Section 1.5.

The Bern carbon cycle model (Siegenthaler & Joos 1992; Joos *et al.* 1996; Fuglestvedt *et al.* 2010; Joos *et al.* 2013) emulates the long-term response of the carbon cycle using

multiple exponentials (i.e advective terms) to reproduce this phenomenon, but the model is in effect modelling a diffusive term. A diffusive term looks similar in the short-term to a series of exponential terms, as suggested by Figure 2.3a if we shift the box-diffusion line upwards, for example by increasing the Revelle Buffer Factor in the model from $b_3 = 0.15$ to $b_3 = 0.25$. By including a diffusive term explicitly we are able to reproduce the same phenomenon using only two adjustable parameters.

We wanted this model to be constrained as much by observation as possible, and using a Bern-style carbon cycle with its three exponential decay curves would have required seven parameters when we only have three carbon-cycle constraints as described in Section 2.6.1. This is another reason why we opted for a diffusive carbon cycle.

The diffusion term has the same form as the simple energy balance model (EBM) used in Equation 2.17, and the mathematical form of the diffusion term is discussed in more detail in Section 2.4.5.

2.3.4.1 Limits to this approach

Modelling the carbon cycle with a single advective term and a single diffusive term is obviously a huge simplification. In reality the carbon cycle is an incredibly complex system involving phenomenon on spatial scales from stomatal evaporation to detritus decomposition to eddy mixing of CO_2 in the mixed layer to regional drying or wetting (Ciais & Sabine 2013).

By modelling the carbon cycle as a single pair of advective and diffusive terms, we are assuming it will behave entirely linearly. In reality the carbon cycle is far more complex than this, and includes small non-linearities that drive the variability of the carbon cycle, and also possibly large non-linearities as will be discussed in the next section (Ciais & Sabine 2013).

While the heavily simplified approach taken in the simple models broadly reproduce the trends of more complex coupled-carbon-cycle-climate models, the temporal variability exhibited in the complex models is not reproduced by the simple models used in this thesis

(Joos *et al.* 2013). In particular, inter-annual variability exhibited by the real world is large and the year-to-year variations in the carbon cycle are not predicted well even in state-of-the-art carbon cycle models (Eby *et al.* 2013), let alone in models as crude as the ones presented in this thesis. This temporal variability comes from the myriad processes which are not represented here.

If we wanted to include some of this complexity, we could begin by adding more processes to the box-diffusion model, though we would need to be aware that we would cease to be able to contain the model with observations and pseudo-observations as outlined in Section 2.6.1. For example we could include more land surface process, an ocean thermocline, or a northern and southern hemisphere to reproduce the seasonal cycles in atmospheric CO_2 concentration.

To include more carbon cycle complexity, we could use a model such as Integrated Model Of Global Effects of climatic aNomalies (IMOGEN), which is still a heavily simplified predominantly land surface carbon cycle model with only a couple dozen equations, but represents far more of the processes in the carbon cycle than are represented here.

The advective and diffusive terms also fail to incorporate a number of processes which could amplify or suppress the increase of atmospheric CO_2 concentrations as temperatures increase, however we have included an emissions-temperature feedback term to represent these carbon cycle feedbacks, which I will discuss next.

2.3.5 Carbon cycle feedbacks

In this model we use a carbon cycle feedback term in Equation 2.7. This is motivated by Cox *et al.* (2000), which suggests that carbon cycle feedbacks could provide a significant positive feedback in the climate system, accelerating climate change. As the temperature increases, the capacity of the biosphere and upper ocean to absorb CO_2 decreases (Friedlingstein *et al.* 2006).

In our model, we represent the climate-carbon cycle feedback in Equation 2.7 by adding an extra, linear temperature-dependent component to the total anthropogenic emissions

emitted each year, E_a. Many of the phenomenon that contribute to this term are an non-linear, such as soil respiration (McGuire *et al.* 1992; Lloyd & Taylor 1994) and wildfires (Page *et al.* 2002; Langenfelds *et al.* 2002). However, when these different phenomenon are aggregated, models have suggested this feedback has been approximately linear since the industrial revolution (Friedlingstein *et al.* 2006; Matthews *et al.* 2009). We assume the time-scale over which sources and sinks of carbon in the biosphere re-equilibrate is of the order one century following Allen *et al.* (2009c) (Prentice *et al.* 2001). Further discussion of the linearisation of this term can be found in Section 2.3.5.2, but first I will discuss a possible alternative formulation of this term.

2.3.5.1 An alternative parameterisation of carbon cycle feedbacks

Allen *et al.* (2009c) and Bowerman *et al.* (2011) use a different carbon cycle feedback term from the one in Equation 2.7. These papers use the formulation

$$E(t) = E_a(t) + b_5 \left(T(t) - \int_0^t T(t') e^{\frac{t'-t}{100}} dt' \right), \tag{2.16}$$

where all symbols are the same as in Equation 2.7. In other words, it was the same as Equation 2.7, but behaved symmetrically for falling and rising global temperatures. Henceforth I will refer to this as the *symmetrical* formulation of the carbon cycle feedback term. As introduced in Section 2.3 I will refer to the alternative formulation introduced in Equation 2.7 as the *asymmetric* formulation.

The difference between these two formulations of the carbon cycle feedback term is shown in Figure 2.3b, where the symmetric lines are dotted lines, and the asymmetric lines are dashed and solid. Because the asymmetric case does not allow a negative carbon cycle feedback, we see that the the blue dotted line does not ever fall, while the blue solid line representing the symmetric formulation does fall beyond 330 years after the emissions pulse. The difference between the responses to a pulse of emissions of the two formulations of the carbon cycle feedback term are explored in more detail in Figures 3.5 and 3.6.

The reason I introduced the asymmetric formulation, is that the symmetric formulation is unlikely to function well when temperatures are falling rapidly due to non-linearities in the carbon cycle (Jones *et al.* 2009; Boucher *et al.* 2012). This term takes carbon out of the atmosphere very quickly when T is falling, leading to unrealistic projections for how the climate will behave once temperatures peak. Since the box-diffusion model was originally designed, a number of papers have come out showing that once emissions fall to zero, temperatures stabilise and CO_2 concentrations fall slowly (Friedlingstein 2005; Hare & Meinshausen 2006; Weaver *et al.* 2007; Matthews & Caldeira 2008; Plattner *et al.* 2008; Eby *et al.* 2009; Lowe *et al.* 2009; Solomon *et al.* 2009; Frölicher & Joos 2010; Zickfeld *et al.* 2010). This is not the behaviour exhibited by the box-diffusion model, or indeed by many of the older simple climate models that are used in the literature, (Meinshausen *et al.* 2011a, b; Vuuren *et al.* 2009; Siegenthaler & Joos 1992; Joos *et al.* 1996) however this term allows us to mimic the behaviour of these more complex models better before temperatures peak, even if the behaviour after the peak is unrealistic. Because the model is tuned to historical observations and C^4MIP, it has only been tuned to periods of rising temperatures, and so the behaviour of the model after temperatures peak is poorly constrained. The behaviour of the model under falling temperatures is discussed in more detail in Section 3.5.

In this thesis I have improved this post-peak behaviour by setting $E = E_a$ in Equation 2.7 when the carbon cycle feedback term would otherwise be removing carbon from the climate system due to falling global temperatures. I do not have a good physical justification for this change given that we do not really know how the carbon cycle system will behave in a regime of falling carbon concentrations and temperatures given that we've never been able to observe the system in this regime. The most likely reason why this could be a good parameterisation of this term would be because of non-linearities in the carbon cycle of the type highlighted by Updegraff *et al.* (1998), Jones *et al.* (2009) and Boucher *et al.* (2012). As we will see in Chapter 3, the asymmetric formulation also does a better job of simulating

historic CO_2 concentrations and replicating the behaviour of more complex models.

Despite the asymmetric carbon cycle feedback term's potential improvements over the symmetric formulation, the asymmetric formation is still extremely simplistic and comes with plenty of caveats, as I will discuss now.

2.3.5.2 Limits to this approach

This carbon cycle feedback term is designed to incorporate a wide number of different phenomenon, none of which are represented explicitly. Some of these processes increase net CO_2 emissions, such as soil respiration (Lloyd & Taylor 1994), while others decrease net CO_2 emissions such as the CO_2 fertilisation of photosynthesis (Friedlingstein *et al.* 2006). Some of these phenomenon can be crudely modelled as being functions of temperature T, while others, such as photosynthesis fertilisation, are more dependant on C (Friedlingstein *et al.* 2006).

This model linearises all of these phenomenon into a single T-dependent term, however this linearization is unlikely to hold for temperatures greater than 3-4°C above pre-industrial temperatures as more feedbacks become activated (Friedlingstein *et al.* 2006; Lenton *et al.* 2008) . The exact amount of warming required to trigger these feedbacks is unknown, but feedbacks such as the hypothesised 'Amazonian dieback', 'Boreal forest dieback', or 'Sahara greening' could cause the linearised carbon cycle feedback term to potentially become unphysical (Lenton *et al.* 2008; Collins *et al.* 2013). Potential 'tipping points' in the climate system are discussed in more detail in Section 3.1.2.

Throughout my thesis I have attempted to only use the model in the regime in which temperatures are increasing or stabilising up to 3-4°C above pre-industrial. In doing so I hoped to keep the model in an operating space that is similar to the range of constraints that were used to tune the model.

We have made this carbon cycle feedback function solely dependent on T and not on C because T is a function of C anyway. However, I could plausibly include two carbon

cycle feedback terms, one dependent on C and one dependent on T which may improve the accuracy of the model.

I could also remove the assumption that this feedback will be linear, and use a different functional form of this term tuned to the output from C^{4}MIP (Friedlingstein *et al.* 2006), or Coupled Model Intercomparison Project Phase 5 (CMIP5), the output from the same suite of models that has been prepared for the IPCC Fifth Assessment Report (AR5) (Taylor *et al.* 2012). I could go further and include complex functions of T, C, t, and possibly cross-terms incorporating more than one of these variables. However, because this term is representing a number of different process which all have different non-linear functional dependences on temperature and other local and global variables, it is unclear which functional form to use in this term if we were to remove the assumption that its behaviour was linear. Until we are modelling the regime in which the carbon cycle feedback has large non-linearities, a simple linear T-dependency is the most basic way of approximating carbon cycle feedbacks, capturing the majority of the effect (Friedlingstein *et al.* 2006).

A more realistic way of representing these processes would be to include them explicitly using an EMIC or ESM, or even possibly an intermediate-complexity model similar to Lenton (2000) or Lenton *et al.* (2008). The worry with using oversimplified models with many free parameters such as Lenton (2000) to represent this sort of phenomenon is that because we do not have good observations of the phenomenon we are attempting to model in the far future, I risk ending up with a model with so many uncertain parameters that it can exhibit a wide range of behaviours. A model of this form that is not well constrained by observations risks being not very useful because we do not know how much of its behaviour is realistic and how much is simply a result of unphysical representations of the processes in the model. Even the box-diffusion model, which was designed to be well-constrained by observations, is poorly-constrained and hence potentially unreliable once temperatures stabilise because we have no observations in that regime.

It is precisely because of the lack of observations of these possible future feedback

mechanisms that they are often not included in models. Advocates of including these feedbacks argue that without them we are underestimating our uncertainty, and particularly underestimating the risk of higher levels of climate change where the dangers and damages are highest (Cox *et al.* 2000; Lenton *et al.* 2008; Lenton 2011).

This term only represents carbon cycle processes, however there are also non-carbon cycle feedbacks that are not represented in the box-diffusion model. For example methane hydrates and permafrost melt have all not been included explicitly this model, which could be included in the future if we wished, particularly if we wanted to make the model more accurate above the 3-4°C point at which I currently stop trusting the results.

Since Cox *et al.* (2000) was published more estimates have been made of the impact of carbon cycle feedbacks, and the Cox *et al.* (2000) estimate appears to have been at the top end of the range of estimates of the magnitude of carbon cycle feedbacks (Friedlingstein *et al.* 2006; Arneth *et al.* 2010; Collins *et al.* 2011; Arora *et al.* 2013). However, this feedback parameter in the box-diffusion model was tuned to the output of C^4MIP and so represents the wide range of uncertainty inherent in that suite of models.

Nonetheless, the carbon cycle feedback term in the box-diffusion does appear to be larger than in many other models, as we shall see in Chapter 3. I will explore the behaviour of the box-diffusion and the implications of a strong carbon cycle feedback in Chapter 3, but before I do that I need to introduce the remainder of the box-diffusion model, and in particular its EBM component.

2.4 A simple energy balance model

The Earth's climate is heavily influenced by the flows and stores of energy in the climate system. The atmosphere plays a key role in maintaining our climate by trapping energy and allowing the effective emissive temperature of the Earth to be different from the temperature at the Earth's surface. From simple energy balance concerns we can show that the Earth's surface would be at 255K (-18°C) if it were a perfectly black body of uniform

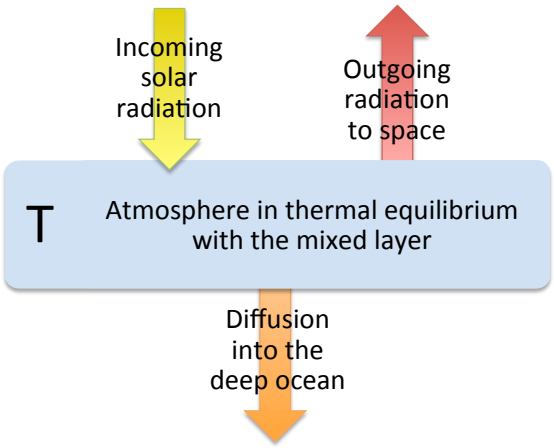


Figure 2.6: A schematic of the different components of the simple EBM, representing the components of Equation 2.17.

temperature without an atmosphere (Taylor 2005). Instead the average temperature on the Earth's surface is about 287K (14°C), where the difference is caused by the atmosphere.

We attempt to model global energy flows using a very simple EBM. We consider only four components: (1) an additional energy flow into the atmosphere caused by CO_2 forcing, (2) additional energy flow out of the atmosphere caused by extra radiation to space due to the increased temperature, which is also affected by feedback mechanisms, (3) additional net energy flow into the deep ocean, caused by a steeper temperature gradient into ocean, and (4) the effective global heat capacity of the atmosphere-ocean mixed layer system.

We represent these components mathematically as

$$a_1 \frac{dT}{dt} = a_3 \ln\left(\frac{C}{C_0}\right) - a_0 T - a_2 \int_0^t \frac{dT(t')}{dt'} \frac{dt'}{\sqrt{t-t'}}. \tag{2.17}$$

a_1 is a fixed heat capacity which we approximate as the effective heat capacity per unit area of a 75m ocean mixed layer. a_3 corresponds to a doubling of atmospheric CO_2 levels causing a forcing of 3.74 Wm^{-2} Ramaswamy *et al.* (2001). a_0 and a_2 are both able to vary, and control the climate sensitivity, and heat flux through the thermocline, respectively. T is the global mean surface temperature anomaly, which is assumed to be linear in the mean mixed layer temperature anomaly. I will discuss each of these terms in turn in the following sections.

2.4.1 Effective heat capacity

The term on the left hand side of Equation 2.17 is the effective heat capacity of the atmosphere and mixed layer. It determines the rate at which the Earth surface warms up for a given rate of net energy absorption by the Earth's surface, measured in units of $\frac{\mathrm{W/m^2}}{\mathrm{K/year}}$. A larger effective global heat capacity results in a larger climate response timescale.

We take the heat capacity of a 75m ocean mixed layer covering 70% of the Earth's surface to be our effective global average heat capacity motivated in part by considering the portion of the Earth that contributes to its effective heat capacity but is not deep ocean, and in part from tuning the model to the results from EMICs (Allen *et al.* 2009c).

2.4.1.1 Limits to this approach

The Earth's surface does not have a single effective heat capacity. An additional layer of complexity we could add to the heat capacity would be to model the land and ocean separately. Land has a relatively low effective heat capacity and so heats up quickly, while the ocean has a much larger effective heat capacity and warms slower. Joshi *et al.* (2007); and Dommenget (2009) have suggested that continental warming is largely caused by simultaneous ocean warming, and not by anthropogenic forcing over the land. Adding the land-ocean warming contrast would potentially result in an improvement in the models' accuracy.

Effective heat capacity is not homogeneous over the ocean, nor over land. For example the effective heat capacity of a stormy well-mixed ocean is larger than that of a stiller stratified ocean (Kawai & Wada 2007). Thus the effective heat capacity depends on non-homogeneous variables such as wind speed, which in turn depends on atmospheric dynamics among other things.

The effective heat capacity in the box-diffusion model is assumed to be a fixed constant. In reality, the Earth's effective heat capacity will change as the planet warms. For example, as the planet warms, changes in ocean circulation could occur and the oceans could become

more stratified (because of larger temperature gradients between the surface and the deep ocean), resulting in less mixing and hence a smaller effective heat capacity (because heat will propagate through the ocean less easily) (Capotondi *et al.* 2012).

2.4.2 Radiative forcing

Radiative forcing (radiative forcing (RF)) is defined as the net increase in downwelling radiation at the top of the atmosphere (Baede 2007). The first term on the right-hand-side of Equation 2.17 is the RF term, and can be replaced with any RF. For example, if we wanted to model the effect of the solar cycle on the climate we could replace the logarithmic CO_2 term with a cyclical function of t representing the change in forcing due to the solar cycle. RFs can be positive or negative as illustrated in Figure 2.7, which shows the contributions to global average RF in 2005 relative to the pre-industrial era. Throughout this thesis I will occasionally shorten RF to simply *forcing*.

The forcing due to CO_2 can be approximated as being logarithmic with increasing CO_2 concentration (Ramaswamy *et al.* 2001). This can be shown using detailed models of radiative transfer that are solved numerically, or for a simpler explanation Asrar & Hurrell (2012) consider how the effective emitting surface of the Earth moves upwards with higher CO_2 concentrations and derive a term that is logarithmic with increasing CO_2 concentration. Throughout this thesis I will be approximating the forcing due to CO_2 as being logarithmic with increasing CO_2 concentration.

2.4.2.1 Limits to this approach

In Figure 2.7 we see that the fractional error on the forcing due to the long-lived greenhouse gases (LLGHGs) is smaller than the fractional error on many of the other sources of RF. However despite the smaller fractional uncertainty on the RF of CO_2, there are still small improvements that we can make to the global mean forcing due to CO_2. In the

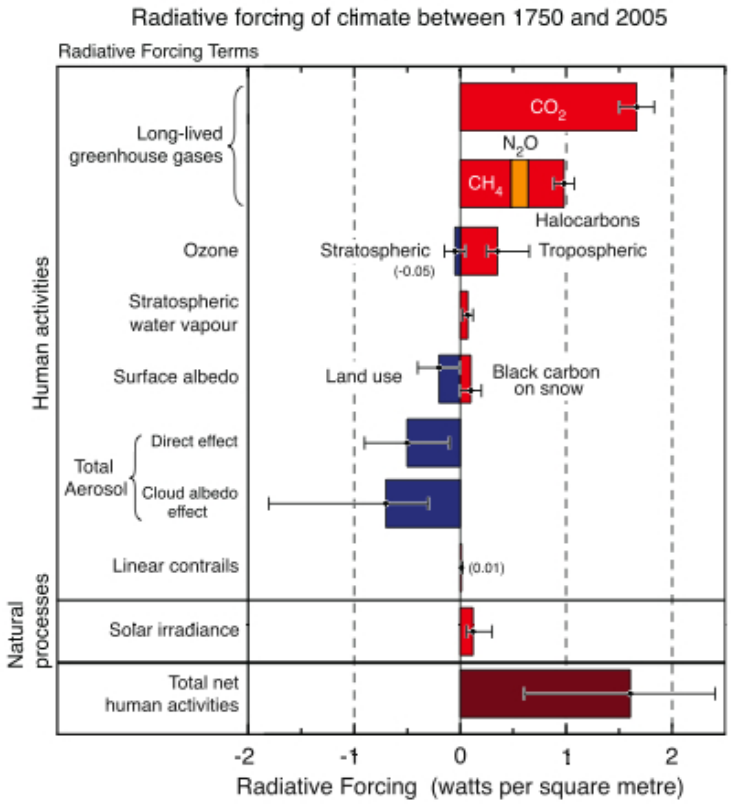


Figure 2.7: Global average RF estimates and ranges in 2005 for anthropogenic GHGs and other forcing agents and processes. Source: IPCC (IPCC 2007)

box-diffusion model, we take the radiative forcing due to CO_2

$$\Delta F_1 = \alpha \ln\left(\frac{C}{C_0}\right) \tag{2.18}$$

where $\alpha = 5.35$ Wm^{-2} from Ramaswamy *et al.* (2001). In Ramaswamy *et al.* (2001) the authors propose this can be made more accurate by modelling the RF due to CO_2 as

$$\Delta F_2 = \alpha \ln\left(\frac{C}{C_0}\right) + \beta\left(\sqrt{C} - \sqrt{C_0}\right) \tag{2.19}$$

where $\alpha = 4.84$ and $\beta = 0.0906$. The authors go further to suggest an even more accurate representation of the RF due to CO_2 of the form

$$\Delta F_3 = \alpha\left(g(C) - g(C_0)\right) \tag{2.20}$$

where

$$g(C) = \ln\left(1 + 1.2C + 0.005C^2 + 1.4 \times 10^{-6}C^3\right) \tag{2.21}$$

and $\alpha = 3.35$.

$\Delta F_1 = 1.81$ Wm^{-2} is 3% smaller than $\Delta F_3 = 1.87$ Wm^{-2} for year 2011 concentrations (assuming 390 ppmv (Ciais & Sabine 2013)), and $\Delta F_2 = 1.92$ Wm^{-2} is 2% larger than $\Delta F_3 = 1.87$ Wm^{-2} again for year 2011 concentrations. These RFs are all within the uncertainty range of the more recent IPCC estimate for 2011 forcing of 1.82 ± 0.18 Wm^{-2} (Myhre & Shindell 2013).

These formulae are obtained from fits to the results of line-by-line radiative transfer models (Ramaswamy *et al.* 2001). These models simulate the propagation of radiation through the atmosphere at different wavelengths, and simulate the interaction of the radiation with the different species of gas at different locations in the atmosphere. The models allow us to estimate the change in net downwards propagation of radiation at the top of the troposphere relative to pre-industrial, which is what we define as radiative forcing. A more complex model, such as a GCM, would contain a radiation call, which would do these line by line calculations to estimate the RF usually every few time steps (Stocker 2011).

So far we have only discussed the RF due to CO_2. As depicted in Figure 2.7, there are many other types of forcing that effect the climate. In chapters 2–5 I am concerned only with CO_2-induced warming. In effect I am approximating the net non-CO_2 forcing to be zero. This approximation happens to hold at the moment, as illustrated in Figure 2.7 where the forcing due to CO_2 is approximately equal to the total net anthropogenic forcing. This approximation is less likely to hold in the future (Meinshausen *et al.* 2011c), which is why in Chapter 6 I build non-CO_2 GHGs into the box-diffusion model.

CO_2 is a well-mixed GHG, meaning that its atmospheric concentration is approximately the same at a given time throughout the atmosphere (Stocker 2011). Shorter-lived radiative species such as black carbon are not well-mixed, and hence their RF is not homogeneous (Shindell *et al.* 2012). I model black carbon in Chapter 6 by assuming it is homogenous, but to incorporate its non-homogeneity into the model I would need to incorporate at least one spatial dimension into the model.

The box-diffusion model considers only annual-mean quantities because it has a time-step of one year. In reality CO_2 concentration has a seasonal dependence, caused in part by the discrepancy between the amount of land in the Northern hemisphere vs the Southern hemisphere (Etheridge *et al.* 1998a). To incorporate this we would need a model with shorter time steps, at least one spatial dimension, and a more complex carbon cycle. This would results in a more accurate carbon cycle model, and in turn a more accurate RF due to CO_2.

2.4.3 Linear feedback term

The second term on the right hand side of Equation 2.17 controls how much radiation in the climate system escapes to space. Feedbacks in the climate system are an important contributor to this quantity. Feedbacks in the climate system are labelled the opposite way from those in engineering; in climate physics a positive feedback is one that makes the climate warmer, pushing the system away from the pre-industrial equilibrium, and negative

feedbacks make the climate cooler, pushing the system towards the pre-industrial equilibrium. Examples of feedbacks in the climate system include water vapour feedbacks, surface albedo changes and changes in cloud cover and composition (Colman & McAvaney 2011).

One way of quantifying feedbacks in the climate system is via the climate sensitivity, ν, which is defined as the amount of warming that would occur if the climate system came into equilibrium following a doubling of CO_2 concentrations relative to pre-industrial (Planton 2013). By considering our model in equilibrium such that $\frac{dT}{dt} = 0$ we can easily show that the climate sensitivity

$$\nu = \frac{a_3}{a_0} \ln(2), \tag{2.22}$$

where ν has units of temperature.

This term is not governed by a single obvious process, but is rather a linearised amalgamation of everything that changes the top of atmosphere (TOA) energy flux imbalance in proportion with temperature. These processes include an array of phenomenon from Boltzmann radiation to cloud and water vapour feedbacks. An alternative way of viewing the first two terms on the right hand side of Equation 2.17 is as the linearised change in the TOA energy flux imbalance with respect to temperature and with respect to the linearised change in the logarithm of the fractional increase in CO_2 concentration, including all atmospheric climate feedbacks.

2.4.3.1 Limits to this approach

In the box-diffusion model we assume that climate sensitivity is approximately constant in the region we assume to be approximately linear below 3–4°C of warming. Above this region however, there may be non-linearities in the climate system which could potentially perturb the climate sensitivity from its current value (Hansen *et al.* 2008; Collins *et al.* 2013).

Even if we ignore some of the more exotic feedbacks in the climate system such as methane clathrates, this term still assumes a linearity which may not exist at higher temperatures. For example, for a homogeneous blackbody with a uniform temperature and

no atmosphere, this term would be a linearisation of the Boltzmann black body radiation curve. This curve is proportional to T^4 and so even between pre-industrial conditions and two degrees of warming, we would expect the gradient of the curve to change by a factor of 2%, and hence not be strictly linear (Blundell & Blundell 2006). For processes such as cloud feedbacks these non-linearites are likely to be more pronounced (Boucher & Randall 2013). Nonetheless, feedbacks are mostly linear in GCMs, and so I would expect the formulation of the model as outlined here to give reasonable results up until the 3–4°C of warming range (Matthews *et al.* 2009; Collins *et al.* 2013).

The box-diffusion model explicitly models very little of the climate system. There is are hundreds or even thousands of process that are represented in more complex climate models, but are not represented here (Stocker 2011). Many of the atmospheric processes are subsumed into this linear feedback term. By representing some of these processes more explicitly, from dynamics to cloud processes, I would expect to have a more accurate representation of the climate, and in particular of variability, but only if we have data to constrain these extra terms with.

To capture more of the complexity in the climate system that impacts on this term, we would need to incorporate the non-homogeneities in this term on Earth. At the most basic level, cloudless regions near the equator emit more outgoing long-wave radiation back to space than cloudy regions or regions near the poles (Taylor 2005). Incorporating some of these inhomogeneity into this model would help it make more accurate projections, particularly of variability.

2.4.4 Efficacies

Though I am only concerned with CO_2 in this chapter, I will discuss efficacies here as they relate directly to Equation 2.17. Thus, this section will examine GHGs beyond CO_2. For a more detailed discussion of non-CO_2 GHGs see Chapter 6.

In this model we have assumed that the climate sensitivity is independent of the GHG causing the RF. In the real world this is not true, and so we must introduce the concept of

efficacy.

Just as forcing is a function of the concentration and gas properties of a GHG, the temperature response is a function of the forcing and the efficacy of that forcing.

Let us define efficacy as the dimensionless quantity,

$$E_i = \frac{\nu_i}{\nu_{\mathrm{CO}_2}}, \tag{2.23}$$

where ν_i is the climate sensitivity for a given GHG i, and ν_{CO_2} is the climate sensitivity of CO_2 (Forster *et al.* 2007). Thus we can say,

$$\Delta T_S = \nu_{\mathrm{CO}_2} E_i F_i, \tag{2.24}$$

where ΔT_S is the equilibrium change in global mean surface temperature, and F_i is the forcing due to a given GHG i. This lets us consider $E_i F_i$ as an effective radiative forcing due to a given GHG (Sausen & Schumann 2000; Hansen *et al.* 2005; Lohmann & Lesins 2005). Note that we can not only include effective forcing and ignore efficacy because efficacy also affects the timescale of the temperature response.

Though estimates vary (Lohmann & Lesins 2005; Hansen *et al.* 2005; Wang *et al.* 1991, 1992; Govindasamy *et al.* 2001), the efficacies of non-CO_2 long-lived GHGs are generally thought to be equal to or slightly higher than one.

Efficacies can be non-unity for a number of contributing factors, but in particular because of geographical and spatial variation between the forcings. Nearly all studies surveyed in the IPCC Fourth Assessment Report (AR4) found that high-latitude forcings have a higher efficacy than tropical forcings (Forster *et al.* 2007). Additionally, forcings that affect the upper troposphere are generally found to have smaller efficacies than forcings that affect the surface. IPCC AR4 suggests that excluding cloud-aerosol affects, there is a high confidence in a linear relationship between global mean RF and the global mean surface temperature response (Forster *et al.* 2007).

2.4.5 Diffusive term

The final term in Equation 2.17 represents the transfer of thermal energy into the deep ocean. We model the transfer of heat into the ocean as turbulent diffusion (Taylor 1915; Vallis 2006). In the real world the meridional overturning circulation (MOC) and southern ocean eddies play a bigger role than turbulent diffusion does in determining the rate at which energy flows into the deep ocean (Rhein & Rintoul 2013), however in the zero-dimensional box-diffusion model we can model these processes as diffusion. We can do this because we find that turbulent diffusion does a reasonable job of representing the global mean flow of energy into the deep ocean, particularly once we are able to tune a_2 and the other adjustable parameters in the model.

The full derivation of how we obtain the exact mathematical form of the $\int_0^t \frac{dT(t')}{dt'} \frac{dt'}{\sqrt{t-t'}}$ term is rather involved (Crank 1979), but we can get an idea of the origin of the $\frac{1}{\sqrt{t}}$ term by considering a semi-infinite slab with a no net heat flux boundary condition on the surface. A delta function temperature perturbation at the surface of a semi-infinite diffusing slab will produce a half-Gaussian temperature profile in the dimension perpendicular to the surface of the slab (Crank 1979). This Gaussian temperature profile will spread out as $\sqrt{t}$, and thus to conserve volume its height, representing the temperature at the surface of the slab, must decrease as $\frac{1}{\sqrt{t}}$ (Crank 1979). Thus for a pulse temperature perturbation, the temperature in the mixed layer of a purely diffusive ocean would decrease as $\frac{1}{\sqrt{t}}$, motivating the form of the diffusive term in Equation 2.17.

2.4.5.1 Limits to this approach

Using a single diffusive term to represent the transfer of heat into the deep ocean is going to result in an oversimplified treatment of long-term energy fluxes in the climate system. A more accurate representation of the deep ocean thermal transport would be to use multi-box ocean model, representing the thermocline and perhaps the North Atlantic Sink to simulate the action of the overturning circulation (Meehl *et al.* 2007).

Per unit area, the North Atlantic Sink draws down more heat into the deep ocean than any other area of the ocean (Gregory 2000; Huang 2003). This is why GCM simulations of climate change usually have less warming over the North Atlantic Sink compared with the average warming over the rest of the ocean (Rahmstorf & Ganopolski 1999; Meehl *et al.* 2007). The North Atlantic Sink is not included in the box-diffusion model. A more comprehensive treatment of effective heat capacity could include a dynamic ocean that can represent the different ways in which heat is transferred into the deep ocean. A dynamic ocean like this would also result in the model representing some of the inter annual variability inherent in the climate system.

Better representation of the ocean would also allow us to better represent the marine currents and ocean circulations that bring bodies of seawater at different temperatures into contact with each other and the atmosphere, creating an 'ocean memory' effect in which the ocean displays a path-dependent hysteresis. At the more extreme and unlikely end of the phenomenological spectrum, better modelling of ocean circulations could model the risk of a shutdown of thermohaline circulation (THC), which would have a significant impact on the global mean climate (Vellinga & Wood 2002), though is unlikely to occur in the near future (Collins *et al.* 2013).

This concludes our introduction to the basic components of model. We have explored how the carbon cycle and climate components of the box-diffusion model work, and discussed some of their limitations. I will discuss the limitations of the model in more detail in Section 3.7.1, but first I will introduce the emissions profiles that will be used to force the model.

2.5 Carbon emissions

Throughout this thesis I will be using different emissions profiles to make projections of the future. In this section I will explain how I have obtained the different emissions profiles that I make use of throughout my thesis. An emissions profile is a possible future anthropogenic

emissions time series. Historical emissions have increased over time from near-zero net anthropogenic emissions in the pre-industrial era, to approximately 10 GtC yr^{-1} today (Meinshausen *et al.* 2011c; Boden *et al.* 2013). In the future they could fall if society engages in rapid mitigation (Moss *et al.* 2010), or they could continue to rise far into the future until fossil fuels run out (Meinshausen *et al.* 2009). What will happen to future carbon emissions is one of the biggest unknowns in future projections of climate change, as explained in Section 1.1.

The IPCC has adopted a set of standardised future emissions scenarios called the Representative Concentration Pathways (RCPs) as reviewed by Moss *et al.* (2010), Meinshausen *et al.* (2011c) and van Vuuren *et al.* (2011). The four main RCPs have names based on the RF that they will achieve by the end of this century. Throughout this thesis I will make extensive use of the RCPs as they span the range of possible future emissions, from extreme mitigation in 'RCP3PD', to business as usual in 'RCP8.5' (Moss *et al.* 2010). The "PD" in RCP-3PD stands for "peak and decline". The main four RCPs are plotted in Figure 2.8. When I need more than four emissions pathways then I use a different methodology to obtain the pathways explained below.

2.5.1 Emissions pathways

When I need more than four emissions pathways then I generate them using the algorithm introduced by Allen *et al.* (2009c). This gives the rate of change of future emissions according to the equations below

$$E_a = \begin{cases} H(t) & \text{for } t < t_0 \\ a\mathrm{e}^{bt} & \text{for } t_0 \leq t < t_1 \\ c\mathrm{e}^{dt^2+ft} & \text{for } t_1 \leq t < t_2 \\ g\mathrm{e}^{ht} & \text{for } t \geq t_2 \end{cases} \tag{2.25}$$

where E_a is the carbon emissions in year t, $H(t)$ is historical emissions data, and a, b, c, d, f, g, and h are constants. t_0 is the year at which historical data is replaced by emissions pathways. The parameters b, and g, representing the initial rate of exponential growth and final rate of exponential decline, depend on the specification of the emissions pathway and

are allowed to vary between emissions pathways. t_1, and t_2 are the times of transitions, and also vary between emissions pathways. The remaining constants are determined by the requirement that emissions are continuous everywhere, and that rates of change of emissions are continuous where $t > t_0$.

To create the emissions pathways, we selected a number of values of parameters b, g, t_1, and t_2. Each combination of these parameters represents a different possible emissions pathway. We selected parameter options such that there are 12,750 possible emissions pathways of the type outlined here. We chose the ranges of the parameters to give a range of emissions pathways with cumulative emissions to 2200 between 0.7TtC and 3TtC. We chose the parameters so that the majority of emission pathways have a maximum rate of emissions decline of less than 4% per year, but we also considered a smaller number of pathways that decrease by up to 10% per year. A subset of the results are plotted in black in Figure 2.8. Note that E_a is measured in tonnes of carbon, as opposed to tonnes of CO_2. To convert to tonnes of CO_2 one would simply multiply our emissions and cumulative emissions values by a factor of $\frac{44}{12}$ (Williams & Follows 2011). These are the emissions pathways that are used extensively in Chapter 4. A subset of these emissions scenarios are plotted in grey and dark blue in Figure 2.8.

Chapter 5 modifies the above emissions pathways as explained in the chapter, and as shown in Figure 5.1. In Chapters 6 and 7 we use the RCPs, and adaptations of the RCPs. Note that the RCPs have slightly different historic emissions from the emissions scenarios described above as illustrated in Figure 2.8, however I do not use these two different scenario families in the same experiment (apart from in Figure 4.2 where it does not matter) so this discrepancy does not affect any of my results.

The emissions scenarios and emissions pathways described thus far are all physically plausible. In order to test the limits of the box-diffusion model and the cumulative emissions concept, I have also created some physically implausible emissions scenarios which I will introduce in Chapter 4.

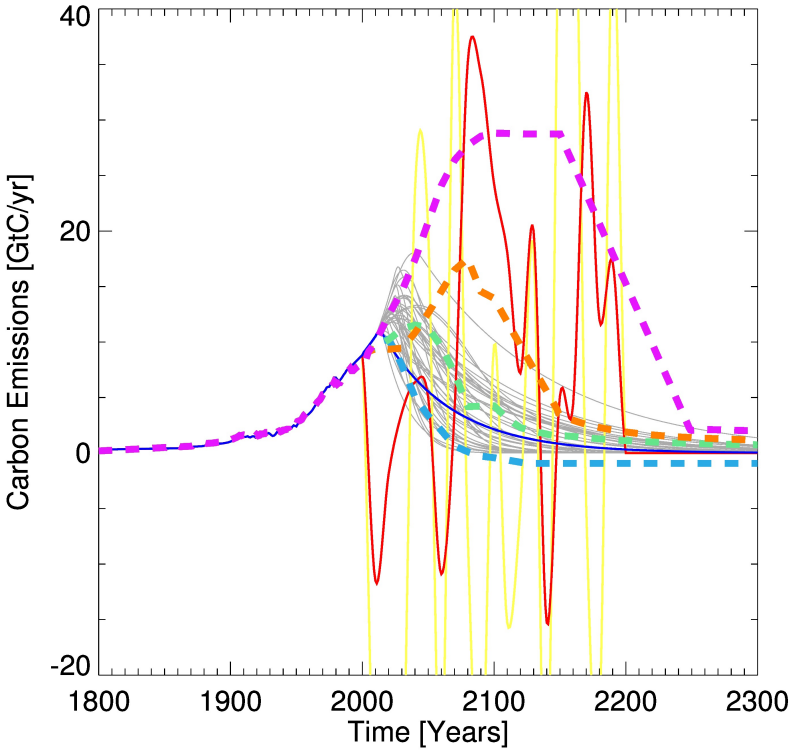


Figure 2.8: Emissions pathways. Pathways in grey are described by Equation 2.25. The dark blue emissions pathway is one example of the grey pathways. The four scenarios in thicker dashed lines, are derived from the Representative Concentration Pathways (RCPs) (Meinshausen *et al.* 2011c). RCP8.5 is in light pink; RCP6 is in orange; RCP4.5 is in light green; and RCP3PD is in light blue. These colours are the same as those used in Figures 3.4 and 4.2. Note that the official RCPs stop in year 2100, so the emissions plotted in colour from year 2100 onwards are technically derived from the unofficial Extended Concentration Pathways (ECPs) (Meinshausen *et al.* 2011c). The red and yellow pathways are an examples of the extreme emissions pathways outlined in Section 4.3. The yellow pathway is highlighted with a yellow cross in Figure 4.2.

2.6 Uncertainty

I introduced the different types of uncertainty that affect projections of climate change in Chapter 1. In Section 3.6 I will consider structural or process uncertainty by comparing two different structural forms of the carbon cycle feedback, but I begin our discussion of uncertainty here by considering parameter uncertainty. I will explain how parameters are chosen in the box-diffusion model, and how we quantify uncertainty. I will begin by discussing the constraints that we use.

2.6.1 Constraints

We need five observations to constrain the five adjustable parameters in the box-diffusion model as outlined in Sections 2.3 and 2.4. In fact we do not have five independent observations so we must instead use the method described below. The five constraints that we use are the:

1. Global effective heat capacity
2. 20^{th} century warming trend attributable to GHGs
3. CO_2 concentration rise attributable to carbon emissions
4. Contribution of the temperature feedback to CO_2 concentration rise
5. Rate of advection of CO_2 in the deep ocean

I will discuss these constrains in turn below.

Global effective heat capacity The global effective heat capacity is defined as the ratio of heat content per unit area to surface temperature rise, where the heat content is dominated by oceanic heat content. We take the best-guess value for ocean heat uptake over the period 1955-1998 from Levitus *et al.* (2005). More recent estimates of upper ocean heat content suggest a slightly higher value (Domingues *et al.* 2008), however it is still well within our

confidence interval. Ben Bronselaer is carrying out work to update this constraint in the box-diffusion model, as discussed in Section 2.7.3. Estimates of surface temperature rise are taken from Brohan *et al.* (2006). Note that the global effective heat capacity is time-scale dependent, as over longer timescales temperature perturbations have more impact on deep ocean temperatures. The global effective heat capacity is different to a_1 in Equation 2.17, because it incorporates both the heat capacity of the mixed layer and the atmosphere, but also the contribution from deep ocean mixing. Hence, the global effective heat capacity is a constraint on both a_1 and a_2 in Equation 2.17.

20^{th} century warming attributable to GHGs This constraint is estimated from the linear temperature trend attributable to GHGs over the 20^{th} century based on "fingerprint" attribution results (Stott & Kettleborough 2002). This trend is then divided into CO_2 and non-CO_2 components via the forcing time series of the relevant GHGs. We use year 2000 RFs of 1.47 and 0.85 Wm^{-2} for the RF due to CO_2 and non-CO_2 GHGs, and assume unit efficacies, which is consistent with the majority of the studies cited in Meehl *et al.* (2007) (Allen *et al.* 2009c). Ben Bronselear is updating this work to use more recent estimates of efficacies, which suggest that non-CO_2 efficacies may be slightly larger than one (Forster *et al.* 2007), as discussed in Section 2.4.4.

CO_2 concentration rise attributable to carbon emissions This constraint can be determined by dividing the observed increase in CO_2 concentrations over a given period by the anthropogenic emissions emitted over the same period (Allen *et al.* 2009c). There are, however, large errors associated with historic anthropogenic emissions data, specifically with estimates of the emissions associated with land-use change (Houghton 2010). We use the land use change emissions from Houghton (2008) and the fractional uncertainty from Denman *et al.* (2007) interpreted as a 5–95% confidence interval. Note that Denman *et al.* (2007) report this range as "likely", meaning a one standard deviation interval; however,

when fit to a log-normal distribution this produces a 5–95% confidence interval that includes negative land-use change emissions, which we believe overstates the uncertainty on this quantity.

Contribution of the temperature feedback to CO_2 concentration rise We do not have two more independent observations to constrain our model. Thus, we use the C^{4}MIP ensemble to contribute our final two constraints. The first, the contribution of the temperature feedback to CO_2 concentration increase, is estimated in the box-diffusion model by by running our model with and without the temperature feedback component of the model activated. Let $C_{A1}(2100)$ be the CO_2 concentration in 2100 under the A2 emissions scenario (IPCC 2007) with the temperature feedback component switched on, and $C'_{A1}(2000)$ be the same but with the temperature feedback component switched off. Our constraint is the quantity $\frac{C_{A1}(2000)-C'_{A1}(2000)}{C'_{A1}(2000)}$, i.e. the fractional contribution of the temperature feedback to CO_2 concentrations. To constrain the model we compare this quantity in our model with the same quantity in the C^{4}MIP models (Friedlingstein *et al.* 2006). One consequence of this methodology is that our simple climate model is no longer independently constrained, and is in fact tuned to the collection of C^{4}MIP models. This means that our model cannot be used to independently confirm C^{4}MIP results, because it will, as a result of this tuning, mimic the basic behaviour of the C^{4}MIP models under the A2 SRES scenario (Nakicenovic *et al.* 2000).

Rate of advection of CO_2 into the deep ocean This constraint is estimated using the results from long-term pulse-injection experiments in EMICs to estimate the rate of advection into the deep ocean (Plattner *et al.* 2008). As with the previous constraint, we use the ensemble mean as the 'best-guess' value for each constraint. We assume the largest fractional deviation from the mean is a two sigma deviation, giving our 5–95% confidence interval.

This concludes my description of the constraints used in the box-diffusion model, and I will now explain how these constraints are used to quantify uncertainty.

2.6.2 Likelihoods

Thus far I have only described how to obtain a single projection of future temperatures for a given emissions profile. In order to establish the errors on this projection, we must go one step further. The adjustable parameters in the box-diffusion model would considerably influence the outcome of our projections if they were perturbed. In order to map the effects of perturbing these parameters, we must run a climate ensemble known as a *perturbed physics ensemble*, which contains a number of *ensemble members*. Each ensemble member uses a different set of values for the five adjustable parameters in the model. In the literature 'ensemble members' are often referred to as 'models' (Tebaldi & Knutti 2007), however I have opted not to use this terminology as it would become ambiguous whether we are referring to an ensemble member or a model such as a GCM.

To find the uncertainty associated with a projection, we need to know the likelihood of the ensemble member that created that projection, and to do this we must carry out a Monte Carlo simulation (Allen *et al.* 2009c; Pawitan 2001). Figure 2.9 contains a schematic of this procedure. For stage (1) consider a parameter space in which each dimension represents a different parameter in our model. Each point in this parameter space represents a set of parameter values and hence a possible ensemble member. In order to obtain our set of inputs for the Monte Carlo simulation we sample throughout parameter space. In this project we use 5000 points creating the 5000 different members of the ensemble. In stage (2) we then numerically solve the coupled differential equations in our model to obtain, in stage (3), an estimate of each of the constraints listed in Section 2.6.1 for each ensemble member.

In stage (4) these modelled estimates of the constraints are compared with the likelihood distribution function of the observed constraints. A log-normal likelihood distribution function, or *likelihood profile*, is associated with each of the constraints in order to represent

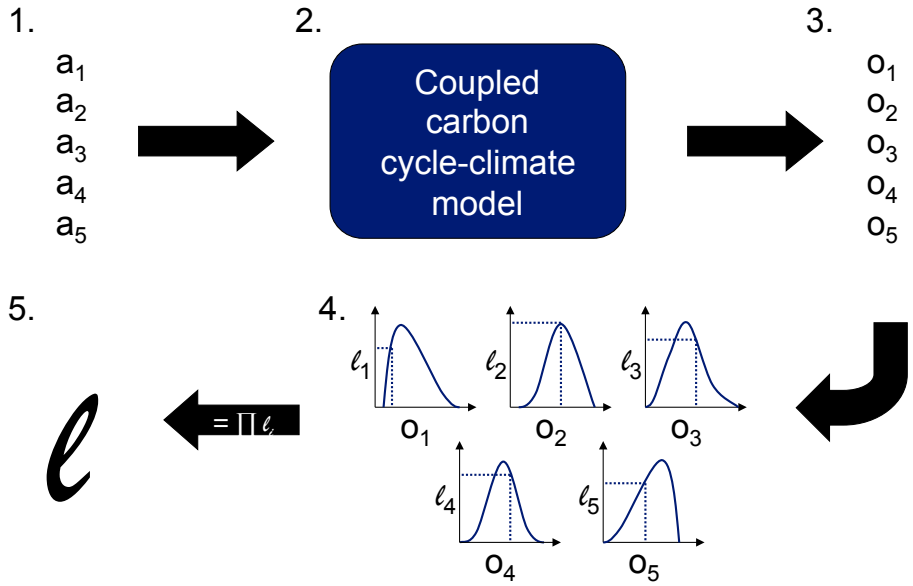


Figure 2.9: A schematic of how the likelihood of different model parameter sets is determined. 1. A set of values of the five model parameters is selected. The method of selection of the parameters is not relevant to determining the parameter set's likelihood. 2. This parameter set is fed into the simple coupled carbon cycle-climate box-diffusion model and 3. a set of model values for five observables is calculated. 4. We also have a set of five observations for the five observables, each with log-normal likelihood distributions taken from the literature.(Allen *et al.* 2009c) The calculated observables are then compared with the real-world observations to give a likelihood for each observation. 5. We take the product of these five observations to determine the overall likelihood of the parameter set presented in 1.

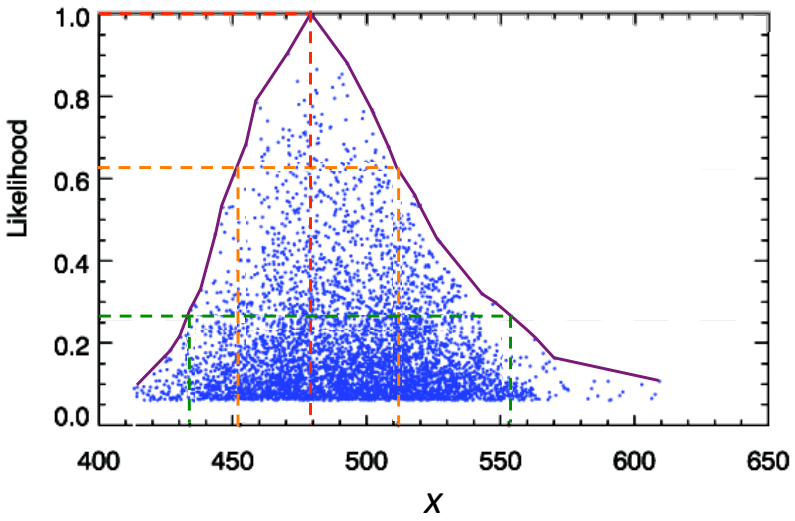


Figure 2.10: A likelihood profile of quantity x. The likelihood profile, in purple, is the outline of all ensemble members, which are plotted in blue. More likely values of x have higher likelihood, where likelihood has been normalised. The IPCC defines x as being *likely* between the yellow lines, and *very likely* between the green lines. Our best guess of x is illustrated by the red line.

its uncertainty. These distributions come from the existing literature and represent the uncertainty on our observations of these parameters in the real world. How these constraints are obtained is discussed in Section 2.6.1, and also by Allen *et al.* (2009c).

Likelihood profiles are traditionally plotted as with variable on the x-axis and likelihood on the y-axis, as illustrated in Figure 2.10. Thus the maximum of the likelihood profile represents the best guess of the variable's value. In stage (5) we assume our constraints are independent, so the likelihood of a given ensemble member or parameter set is taken to be the product of the likelihoods associated with each of the estimates of each of the constrained quantities. These likelihoods are then normalised so that the most likely member of the ensemble has a likelihood of unity.

When choosing parameters in stage (1) of this process we can sample however we please, provided that we sample densely enough to determine where the edge of the likelihood distribution lies.

This likelihood profiling methodology takes a frequentist approach to parameter selection (Pawitan 2001). We do this as part of an effort to make no prior assumption about the

values of the parameters and follow only the data. We are motivated here by Frame *et al.* (2005) and their finding that the probability distribution of climate sensitivity is heavily influenced by the choice of prior, and that there is no obviously correct prior to choose. By using a likelihood method we avoid the use of expert elicitation or subjective priors, and use instead only observations and pseudo-observations from model output, resulting in a more objective quantification of uncertainty (Pawitan 2001). Note that as a Bayesian, I am not claiming here that this approach is more likely to accurately represent future events, merely that estimates obtained using this approach are obtained more objectively.

By using the likelihood methodology we can sample our parameters from anywhere in the distribution using any sampling algorithm that we like and we will get the same likelihood profile, provided that we have sampled densely enough near the peak of the likelihood curve to find its edges. The sampling method we use attempts to do this, however ideally it would sample significantly more points near the edge of the likelihood profile, particularly at higher likelihoods as there are only a few points in the top fifth of the profile. The random forests method of sampling could help us achieve this more effectively (Leo 2001).

While this likelihood-based approach does have its advantages it also has significant drawbacks; for example, we are unable to assign a probability distribution to any of our results without first choosing a prior. Because we only consider likelihoods we are unable to talk about the chance of a parameter value lying within a given range. In Figure 2.10 for example, we cannot talk about the chance of x lying between 444 and 478. We can only say that the likelihood of the observations is halved in the vicinity of $x = 444$ compared with the likelihood of the observations at $x = 478$, as determined from the likelihood profile.

In order to obtain confidence intervals we consider how well different models fit the observations. In this framework the likelihood of a model is independent of the variable we are measuring in that model, so we can use this simplification to estimate the confidence interval on any output of the model. We simply consider all of the models which fit the

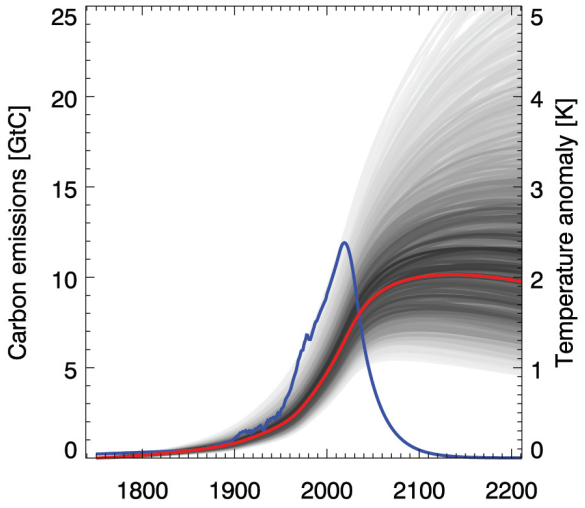


Figure 2.11: Uncertainty in the temperature response to an emissions pathway. The emissions pathway is made up of only CO_2 emissions with a cumulative total of 1TtC. The emission scenario is plotted in blue and uses the left axis. The shaded plume, which is read off of the right axis, is the temperature response to the emissions, with darker curves representing more likely outcomes. The red curve, which can be read off of the right axis, represents the most likely temperature response. Source: My MPhys thesis on "modelling the behaviour of the coupled climate-carbon cycle system".

observations within the y_1 to y_2 confidence interval, and calculate the range they span for a given variable, x. We take this range in x to be the y_1 to y_2 confidence interval, because these models to fit the observations within the y_1 to y_2 confidence interval. The 67% confidence interval is termed by the IPCC as *likely* and is shown in yellow in Figure 2.10. The 90% confidence interval is termed by the IPCC as *very likely* and is shown in green in Figure 2.10.

Likelihood profiles are often plotted as in Figure 2.10 (Pawitan 2001), but they can also be plotted as likelihood plumes as in Figure 2.11. Here more likely ensemble members are plotted darker and less likely ensemble members are plotted lighter. We plot the lighter ensemble members first and then the darker ones on top. The red curve shows the most likely ensemble member. The figure shows the temperature response to an idealised emissions pathway as calculated using the box-diffusion model described throughout this chapter.

Throughout this thesis, when I have been unable to show the entire likelihood profile or plume, I have shown only low, best-guess, and high response parameters. The low- and high-response parameters span the 17-83%, IPCC 'likely' confidence interval, and are

selected as the highest and lowest response models with high enough likelihood in the likelihood profile of the variable in question. For example, in Figure 2.10 these would be the left-most and right-most models above the horizontal yellow dotted line, and will be located close to the intersections of the yellow dotted lines. Note that as we are considering the CO_2 response these likelihoods quantify the uncertainty in the carbon cycle and climate response, but not the uncertainty in the non-CO_2 concentrations, as discussed in Section 6.4.

Now that I have described the box-diffusion model that will be used throughout this thesis, in the next chapter I will discuss how I and others have tested the model by comparing it against historical data and future simulations. Only once I have verified that it is a suitable model will I go on to use it to obtain new results in Chapter 3.

2.7 Ongoing and future work

2.7.1 Advection and diffusion

One area of investigation that merits attention is the functional form of heat and carbon transport into the deep ocean in the box-diffusion model. While I have modelled the carbon cycle and the EBM using one advective and one diffusive term, I could also have used two advective terms or for that matter two diffusive terms in each of the two components of the model. The model would need to be re-tuned to observations and C^4MIP each time a term in the model is changed. Ben Bronselaer[1] has taken up some this work, and has a working version of this model that can switch between advective and diffusion terms in the ocean for either heat or carbon transport.

If I were to change all of the diffusion terms in the box-diffusion model for advection terms, I would then have a analytically tractable model that could give explicit equations for concentration and temperature profiles. I started this work during my Masters' project by solving a climate model explicitly that had only one advection term in the carbon cycle

[1]Corresponding address: bronselaer@atm.ox.ac.uk

and no diffusion in either the carbon cycle or EBM. I showed that any emissions scenario that could be made up of trapezia could be modelled explicitly. It would be interesting to extend this work to a model with two advective terms in the carbon cycle and an advective ocean in the EBM. It would also be interesting to consider emissions scenarios that have non-linear terms, for example quadratic or exponential terms, as these could most likely also be solved explicitly.

2.7.2 Non-stationary climate sensitivity

In the box-diffusion model we have assumed a constant climate sensitivity, as discussed in Section 2.4.3. If we wanted to model climates with non-station climate sensitivities, we could make the climate sensitivity temperature or time dependent. We could also potentially add some threshold-dependent behaviour to represent possible *tipping points* (Lenton *et al.* 2008).

2.7.3 Uncertainty

The likelihood profiles used to constrain the box-diffusion model are now several years out of date (Allen *et al.* 2009c). These could be updated to reflect more up-to-date values. This work has been taken up by Ben Bronselaer.

As discussed in Section 2.6.2, it would be useful to improve our parameter sampling algorithm so that it samples more densely around the edge of the likelihood profile, so that the profile is defined by more points. Figure 2.10 shows how few points actually define our likelihood profile after using the current algorithm. We could implement a random forests sampling method to help sample more heavily in the important areas (Leo 2001).

In this thesis I will only use likelihoods, which as explained in Section 2.6.2, are not probabilities. The likelihoods could be transformed into probabilities using the Jeffery's prior (Jewson *et al.* 2010a, b, c), which is thought to be an objective prior. This work has been taken up by Dan Rowlands[2].

[2]Corresponding address: rowlands@atm.ox.ac.uk

2.8 Chapter summary

2.8.1 The box-diffusion model

The box-diffusion model was motivated by a desire to have a simple climate model that was both transparent and accurate.

We wanted a model that was simple enough to be communicated as a handful of equations, rather than as thousands of lines of code. Simple models like MAGICC and Hadley Centre Simple Climate-Carbon Cycle Model (HadSCCCM1), while they represent more of the complexity of the climate system than the box-diffusion model, took two long academic papers to explain (Meinshausen *et al.* 2011a, b). We wanted a set of equations that represented physical phenomenon, that could be written down from memory. Hence we designed the model around a simple set of six equations, Equations 2.3–2.7 and 2.17.

The Bern simple carbon cycle model coupled with a simple EBM without ocean diffusion such as the model used by Shine *et al.* (2005) would have fulfilled this need for simplicity and transparency (Joos *et al.* 1996; Myhre & Shindell 2013). The model used by Shine *et al.* (2005) can be written down as only two equations, but it is over-simplistic. We designed the box-diffusion model in response to criticisms levelled at simple models that they did not have a carbon cycle feedback term (Cox *et al.* 2000; Friedlingstein *et al.* 2001, 2003; Matthews 2006; Jones *et al.* 2006). Indeed, the Shine *et al.* (2005) model has no carbon cycle feedback at all. This is why we did not simply use one of the many pre-existing simple models (Jones *et al.* 2006; Meinshausen *et al.* 2011a, b; Randall *et al.* 2007), and instead created the box-diffusion model which is only a few equations but still includes a carbon cycle feedback term. Additionally, a simple EBM like that used by Shine *et al.* (2005) does not include the transport of heat into the deep ocean, and so misses out on an important heat sink in the system and the associated slower timescales of the response (Solomon *et al.* 2009; Friedlingstein *et al.* 2011; Collins *et al.* 2013).

The main motivation for creating the box-diffusion model was the ability to have the adjustable parameters constrained by observations and psuedo-observations in the form of

two modelling results from C^{4}MIP (Allen *et al.* 2009c; Friedlingstein *et al.* 2006). By using the maximum likelihood methodology, we were able to use the observations and pseudo-observations from the C^{4}MIP ensemble to quantify the uncertainty on the model projections (Pawitan 2001). In order to do this, we needed to limit the number of free parameters to only five. The Bern carbon cycle model alone has seven free parameters, and so we were motived to use a diffusive term to represent some of the shorter-timescale exponentials that are used in the Bern model (Joos *et al.* 1996; Myhre & Shindell 2013). By using a single advective term and a single diffusive term in the three-box carbon cycle in the box-diffusion model, we were able to significantly reduce the number of adjustable parameters in the carbon cycle to only three, which could then be constrained by our observations and modelling results. A major benefit of this method is that we do not have to use expert elicitation or model tuning to get the required response from the model, which come with their associated biases (Kynn 2008; Mauritsen *et al.* 2012). Instead we can use the more objective maximum likelihood methodology and constraints (Pawitan 2001; Allen *et al.* 2009c).

2.8.2 Emissions

In designing the emissions pathways that we have used in Chapters 4 and 5, we wanted the same level of simplicity as we obtained when designing the box-diffusion model. That is why we chose emissions floors that consist of a fixed increasing rate of change, before transitioning linearly to a fixed decreasing rate of change, as described by Equation 2.25. This family of emissions pathways are quite similar to those we used as part of the AVOID project, particularly once I introduce emissions floors in Chapter 5, although ours have a simpler mathematical form and hence have a continuous gradient everywhere (Huntingford *et al.* 2012).

Compared with the equal quantile walk (EQW) emissions pathways developed by Meinshausen *et al.* (2006), our emissions pathways arguably explore a more limited fraction of

possible emissions pathways. In order to explore a wider range of possible emissions pathways, I will also make use of the RCPs throughout this thesis, as they were designed to span the range of possible emissions trajectories from business as usual in RCP-8.5 to extreme mitigation in RCP-3PD (Moss *et al.* 2010; van Vuuren *et al.* 2011). The RCPs have the additional feature of not being smooth emissions curves like our pathways or those we used as part of the AVOID Project (Huntingford *et al.* 2012), as they were constructed in integrated assessment models (IAMs) (van Vuuren *et al.* 2011; Meinshausen *et al.* 2011c). I will also introduce more extreme emissions pathways in Chapter 4 in order to test the box-diffusion model and the concept of cumulative emissions in more detail.

2.8.3 Limitations of the box-diffusion model

The box-diffusion model has many limitations, which I discuss in detail in Sections 2.3 and 2.4, and will discuss further in Chapter 3. Here I will highlight two of the models' largest limitations that I have returned to throughout this chapter, namely the lack of future observations, and the lack of complexity.

The box-diffusion model is called a simple model because it lacks complexity. Apart from the carbon cycle feedbacks term, it is a linear model and suffers from the drawbacks that linear models of the climate system have. It has no dynamics, and hence no inter-annual variability. The model has no ocean dynamics, and so cannot simulate the long-term 'memory' of the ocean that gives rise the the path-dependent behaviour exhibited in GCMs. Another potentially large difference between the behaviour of the box-diffusion compared with the behaviour of the real world, is that the box-diffusion model does not display any 'tipping point' behaviour, or non-linearities associated with higher temperatures.

While a lack of complexity is a drawback of simple models, a lack of constraints on the behaviour of climate models in the far future is a problem facing all climate models. We have never observed the climate under conditions of falling global mean temperatures, and so it is difficult to know how the system will behave. The behaviour of the box-diffusion model in this regime is almost entirely unconstrained by the historical observations and

features of C^{4}MIP that we have used to constrain the model. As such the model's behaviour in this region is an unreliable indicator of what would occur in the real world.

We do have some constraints on the model, in the form of historic observations, and we can also test the box-diffusion model against more complex models. I will use these constraints and more complex model output to test the behaviour of the box-diffusion in the next chapter in order to validate it and explore its limitations further, before moving on to use the model to explore emissions targets in Chapter 4.

Chapter 3

Testing the box-diffusion model

3.1 Testing climate models

A model is only as good as the predictions it makes (Goody *et al.* 1998). Models that make predictions about the near future, about phenomenon that are easy to observe and measure, are the easiest to test (Gleckler *et al.* 2008). One simply takes the prediction, and compares it with the observations of the period that the model was attempting to predict, and finds out whether the model predicted the phenomenon in question accurately or not. For example, this is how numerical weather predictions (NWPs) are tested (Brier 1950). In NWP, a skill score is a measure of how well a model predicts the phenomenon that it is attempting to predict. Skill scores for NWPs have been increasing steadily over the past few decades as a result of better science, better models, more computing power, and more and better observations (Simmons & Hollingsworth 2002).

The problem with testing climate models is that it takes a very long time to test them. Climate can be thought of as the 30-year average weather (Baede 2007). Thus predictions made by climate models such as the ones used in this thesis are only about 30-year averages, and not about the temperature on a given day, or even averaged over a given year. Because of the long timescales involved in these predictions we have to wait a long time before we can verify whether the predictions they make are accurate (Goody *et al.* 1998). As a result we will not be able to properly test the predictions of climate models made today for another 30 years.

3.1.1 Hindcasting

Because one cannot simply test the predictions of the climate models with future observations, one must take a different approach, known as hindcasting (Kim *et al.* 2000; Sotillo *et al.* 2005; Stefanova *et al.* 2011). Hindcasting is making a prediction about the past, and testing it against historic data. Hindcasting can help verify models against historic data, and in particular is useful for testing that climate models are able to reproduce the phenomenon and processes that we expect will be important in determining future climate. For example, we would expect that a GCM that is unable to reproduce the present day El Nino response, or the North Atlantic Oscillation (NAO), would not have as reliable projections of future climate change as a GCM that was able to reproduce these phenomenon (Landsea & Knaff 2000).

We need to be careful when hindcasting however, to ensure that we are not simply testing against the data that we used to create the model. For example, when creating the box-diffusion model we tuned it to the 20^{th} century temperature trend (Allen *et al.* 2009c). Now if we were not able to reproduce the 20th century temperature trend after tuning our model to it we would have a particularly bad model, but even if we were able to reproduce the trend we still wouldn't know whether the model is able to predict anything else because it was tuned to the very thing we tested it against. Thus a good hindcast in one in which a model is able to reproduce phenomenon that the model was not tuned with (Goody *et al.* 1998).

One of the problem with hindcasting is that it can be difficult to ensure that the data that you are testing the model against is completely independent from the data that was used to create the model. If the data used to tune the model is not completely independent from the data you are using the test the model, then it is not truly a test, and risks causing overconfidence in a model's abilities (Goody *et al.* 1998).

Another potential problem with hindcasting is that of overfitting (Kug *et al.* 2007). A model that is tuned to historic data can fit the data too well, such that it ends up overfitting to

variability that it should not really be predicting as the variability is too random or chaotic in nature. This may result in the model hindcasting well, but being unable to produce reliable forecasts. This can be the case with some statistical models that are based on statistical trends in the data rather than physical processes, which is why climate models are based on physical principles rather than statistical trends (Stocker 2011).

However, even a model that produces successful hindcasts cannot be guaranteed to make successful forecasts.

3.1.2 Regime change

One of the reasons why models that produce successful hindcasts may not necessarily produce reliable forecasts is if there is a regime change. This is a way of saying that the model is being used to extrapolate outside of the region in which it is valid (Lenton *et al.* 2008; Alley 2003).

For example, the statistical models that are used by some firms to make predictions about the stock market can fall prey to this regime change phenomenon (Narang 2013). These statistical models can function well for several months, making significant profits, but then a seemingly insignificant change occurs in market dynamics and the black box models go on to lose millions through poor predictions in the new market regime (Narang 2013). These regime changes can be the result of a change in any number of variables or processes, and represent a shift in the system such that the system exhibits behaviour that is no longer simulated by the model.

The reason the box-diffusion model may not be valid above 3-4°C of warming, as discussed in Section 2.3.5.2, is because above this threshold the risk of non-linearities dominating the climate system becomes significant (Fankhauser *et al.* 2001; Meehl *et al.* 2007; Smith *et al.* 2009). In this region of possible regime change, feedback processes that are not included in the box-diffusion model and many other models may become dominant drivers on the climate system. The feedbacks that could have this effect if they were to become particularly active include: melting permafrost releasing frozen methane (Tarnocai

et al. 2009; Dolman *et al.* 2010; Koven *et al.* 2011); the release of methane clathrates from shallow ocean beds (Archer 2007; Fyke & Weaver 2006; Lamarque 2008); Amazonian dieback (Malhi *et al.* 2009; Good *et al.* 2011; Nobre & Borma 2009) and Arctic sea ice loss (Lindsay & Zhang 2005; Eisenman & Wettlaufer 2009) among others (Lenton *et al.* 2008). Each of these phenomenon has the potential to cause the climate to warm faster than expected by either causing an injection of additional methane, CO_2, or by shifting the Earth's albedo. It is unclear how large these effects could be, nor when they could occur (Meehl *et al.* 2007; Eisenman & Wettlaufer 2009; Nobre & Borma 2009; Smith *et al.* 2009). Nonetheless, these phenomenon have the potential to force the climate system in ways that are not represented in the box-diffusion model, which could not be tested using hindcasting, or even by verifying the model up to 2°C of warming.

Having discussed climate models in abstract, I will now start testing the box-diffusion model. I will begin by testing the model against historical data, and then against other projections of future climate change. Later in this chapter I will use concepts taken from the GHG metrics community (Lashof & Ahuja 1990; Tol *et al.* 2012; Joos *et al.* 2013), to compare the behaviour of the box-diffusion model with the behaviour of other carbon cycle and climate models.

3.2 Testing against historical data

In Section 3.1 I discussed how hindcasting can be used to test models. In Figures 3.1 and 3.2 I compare the box-diffusion against historical observations. I will begin by discussing the comparison with historic temperatures.

3.2.1 Reproducing historical temperatures

Simple climate models such as EBMs do a poor job of reproducing the historic temperature because they lack so many important phenomenon that help determine the historical temperature record (Randall *et al.* 2007). For example they lack dynamics, and so are unable

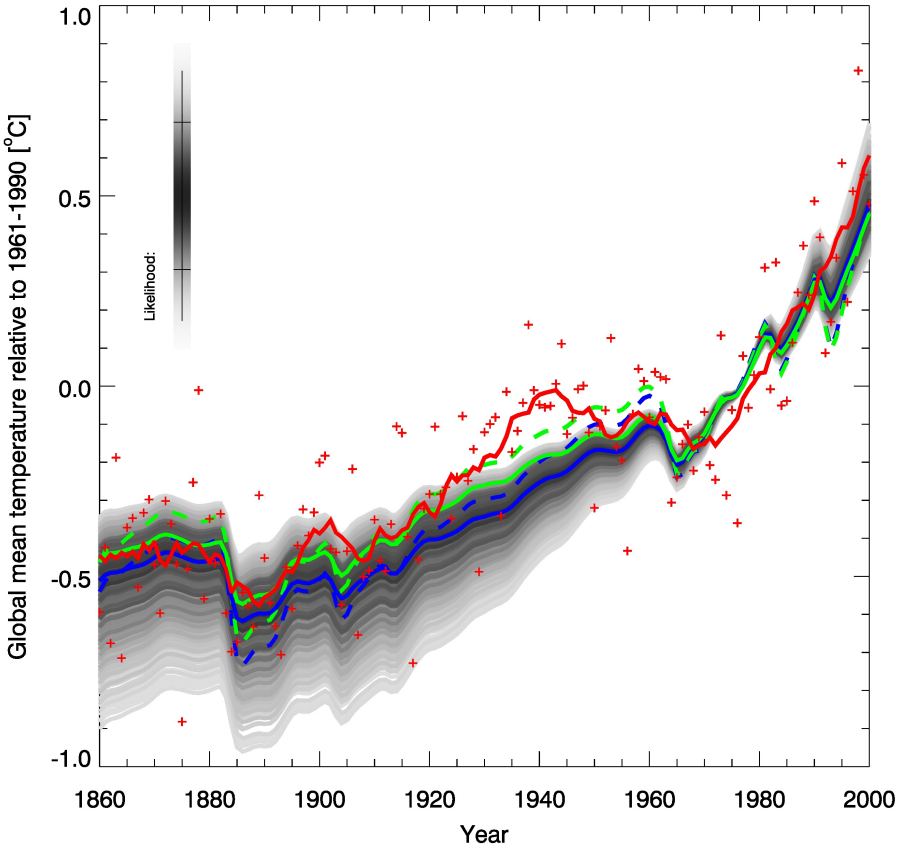


Figure 3.1: Comparing the box-diffusion model against historic observations of surface temperatures from HadCRUT4 (Jones *et al.* 2012; Kennedy *et al.* 2011; Morice *et al.* 2012). Red crosses are annual mean temperatures from HadCRUT4 and the red line is a 10-year running mean of HadCRUT4. The 2σ error on the HadCRUT4 decadally smoothed trend is approximately between 0.05–0.2°C, and is smaller during 1961–1990, and larger near the beginning of the time series, but see Morice *et al.* (2012) for a more detailed uncertainty analysis on the observations. Blue lines are output from the box-diffusion model with emissions as the input, with the carbon cycle turned on, and natural and anthropogenic non-CO_2 forcings taken from the unofficial RCP-20c3m forcings from Meinshausen *et al.* (2011c). Green lines are output from the climate component of the box-diffusion model, but with historic CO_2 concentrations (Keeling *et al.* 2013; Etheridge *et al.* 1998a, 1996) and natural and anthropogenic non-CO_2 historic forcings (Meinshausen *et al.* 2011c) as the input and temperatures as the output. Dashed green and blue lines use the natural and non-anthropogenic forcings from Meinshausen *et al.* (2011c) exactly as specified. Solid blue and green lines use half the volcanic forcing prescribed by Meinshausen *et al.* (2011c) following the result from Stone (2005) that the EBM response to volcanic forcings are too large because of EBM's lack of dynamics. As in Figure 3.2, the grey-shaded plume shows the 5-95% uncertainty on the solid blue line, as shown in the vertical scale bar. The vertical scale bar shows the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively.

to reproduce inter annual variability without a stochastic term of some description (North & Stevens 1998), which the box-diffusion model does not have. As a result we find that the simple model over-reacts to volcanic forcing, as illustrated by the dashed blue line in Figure 3.1. This agrees with the work of Stone (2005), who suggested that EBMs' response to volcanic forcing should be approximately half as large as their response to GHGs. This is why I have included the solid lines in Figure 3.1, which have their volcanic forcing halved.

In Figure 3.1 we find that the EBM is able to reproduce historic temperatures reasonably well, apart from the period around 1940 when temperatures are higher than those simulated by the EBM. This is in keeping with the results of Stone (2005), who was able to fit the temperature trend of the past century with an EBM, but not the period around 1940. Stone also used different climate sensitivities, global heat capacities and ocean diffusivities for each of his four different types of forcing: GHGs, volcanic, solar and sulphate aerosols. Here we use have used the same parameters for each of the forcings, but have halved the volcanic forcing for the dashed curves in Figure 3.1, and so would expect a worse fit than that of Stone (2005). Many GCMs also poorly simulate the warming trend between 1920 and 1940, as it has been hypothesised that it is largely due to natural variability that is not simulated by current GCMs (Delworth & Knutson 2000; Meehl *et al.* 2004; Crook & Forster 2011).

The volcanic forcing used here is from Meinshausen *et al.* (2011a), and is predominantly positive with some large negative spikes, giving a net forcing when averaged over time of close to zero. This is why the dashed full volcanic forcing lines are usually above the solid half volcanic forcings lines, but then occasionally drop far below them when a volcano occurs.

As expected, the model driven with historic CO_2 concentrations (in green in Figure 3.1) reproduces historical temperature trends better than the model driven with historic CO_2 emissions (in blue). Having discussed the ability of the box-diffusion model to represent historical temperature trends, I will now consider its ability to reproduce historical CO_2

concentrations.

3.2.2 Reproducing historical CO_2 concentrations

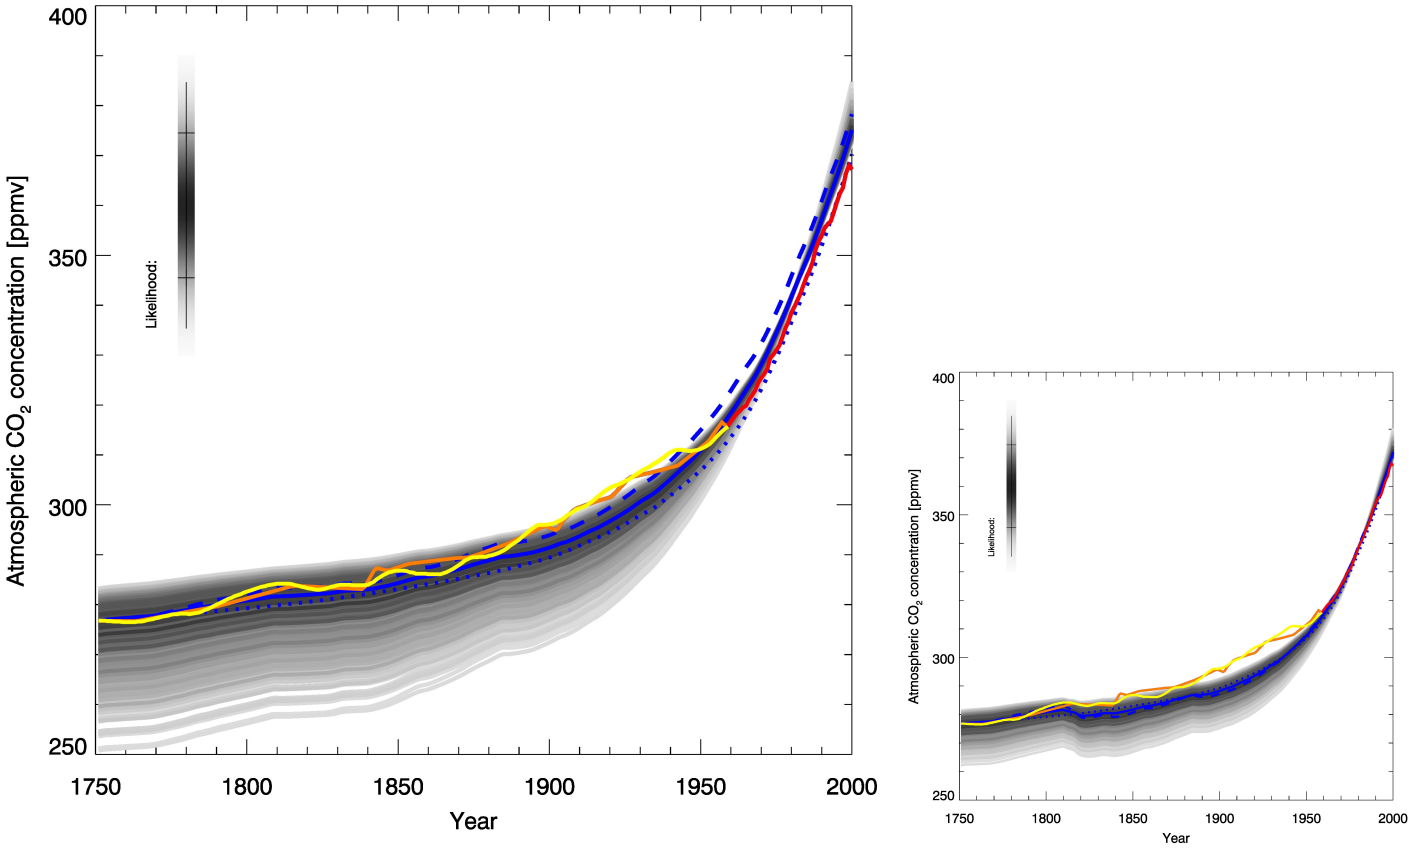
In Figure 3.2 we see a comparison of historic CO_2 concentrations (in yellow, orange and red) with the output from the box-diffusion model (in blue and greyscale).

The box-diffusion model output simulates less likely projections relative to a fixed constant pre-industrial background concentration, as illustrated in Figure 2.11. This output gives almost no uncertainty in pre-industrial concentrations, which in reality still have some uncertainty associated with them (Neftel *et al.* 1994; Etheridge *et al.* 1996, 1998a). On the other hand, the model output presents significant uncertainty in present day concentrations, which are actually rather well known (Keeling *et al.* 2013). To correct this discrepancy, the less likely projections in this plot have been calculated relative to the 1961-1990 average concentrations in the most likely ensemble member so that ensemble members of all relative likelihood have the same 1961-1990 mean concentrations. The most likely ensemble members is unchanged by this transformation. This is why there is almost no uncertainty for atmospheric CO_2 concentrations during the 1961-1990 period.

In making this plot, the box-diffusion model incorporated forcings from non-CO_2 components and included the carbon cycle-climate coupling so as to simulate carbon cycle feedbacks from temperature changes more accurately. Hence, the uncertainty presented here includes both uncertainty in the carbon cycle response, but also in the temperature response via the carbon cycle temperature feedback, as used throughout this thesis and as outlined in Section 2.6.2.

As we would expect, the model reproduces the instrumental record rather well (in red in Figure 3.2), as the carbon cycle in the box-diffusion model is tuned to the historical record, among other constraints. The model has a fixed pre-industrial concentration, and so it is unsurprising that it simulates pre-industrial concentrations well.

The model does less well in the first half of the 20^{th} century where it simulates concentrations a few ppmv lower than the observations taken from from ice cores (in orange and



(a) Historic concentrations with asymmetric carbon cycle feedbacks

(b) Historic concentrations with symmetric carbon cycle feedbacks

Figure 3.2: Comparing the box-diffusion model against historic observations of CO_2 concentrations. Red lines are the instrumental record, taken from CDAIC (Keeling *et al.* 2013), which has an error of less than 1ppmv. Yellow lines are historical observations taken from the Law Dome ice core data (Etheridge *et al.* 1996, 1998a), and have a 1σ error of 1.2ppmv (Etheridge *et al.* 1996). Orange lines are observations taken from the Siple Station ice core data (Neftel *et al.* 1994), with a 1σ error of 5ppmv on the concentration measurement during pre-industrial times, and two years on the measurement of the date throughout the sample. The dashed blue line is output from the box-diffusion model with emissions as the input, with the carbon cycle turned on, and natural and anthropogenic non-CO_2 forcings taken from the unofficial RCP-20c3m forcings from Meinshausen *et al.* (2011c). The solid blue line is the same as the solid blue line, but with historic volcanic forcing from Meinshausen *et al.* (2011c) halved. The dotted blue line is the same, but with no historic volcanic forcing included. The grey-shaded plume shows the 5-95% uncertainty on the solid blue line, as shown in the vertical scale bar. The vertical scale bar shows the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively. Panel (b) is as panel (a), but panel a uses the asymmetric carbon cycle feedback configuration, and panel (b) uses the symmetric configuration.

yellow in Figure 3.2). This is likely due to a number of reasons, but the lack of dynamics in the simple model may be a factor as it leads to poor representation of different forcings, particularly volcanic forcings (Stone 2005). Because this model has a relatively strong carbon cycle feedback as we will discuss in Section 3.3, the over-exaggerated response to volcanoes is causing the emissions feedback term to halt the growth of atmospheric CO_2 concentrations unnecessarily via a temporary zeroing of the emissions feedback term E_a in Equation 2.7. We see this occurring in the model particularly after Krakatau in 1883 (Ammann 2003; Fischer *et al.* 2007), when we compare the behaviour of the dashed blue line, which has full volcanic forcings, against the dotted line, which has no volcanic forcings. It is also likely that some of this discrepancy is due to natural variability, which is not represented in the box-diffusion model, as CO_2 concentrations in this period are also poorly reproduced by the C^4MIP models, with the exception of HadCM3LC (Friedlingstein *et al.* 2006). An additional reason for the poor response is this area is likely due to the discrepancy between the impulse response of the box-diffusion model and the impulse response of more complex models, as I will discuss in Sections 3.4 and 3.5.

One additional feature of Figure 3.2 that is deserving of comment is the fact that the simulations with full volcanic forcing have higher concentrations that the simulations without volcanic forcings. The volcanic forcings provided by Meinshausen *et al.* (2011c) are predominantly positive, with large occasional negative spikes caused by volcanoes giving a time-integrated net radiative forcing which is approximately zero. However the emissions feedback term on the right hand side of Equation 2.7 is defined such that if its output is negative then the term is set to zero. Thus when we have greater variability in temperature trends we will have slightly greater overall CO_2 emissions simulated in this model. This phenomenon is what causes the model to simulate higher carbon concentrations for scenarios with larger volcanic forcings. It is also this phenomenon that produces the difference between the symmetric and asymmetric formulations of the carbon cycle feedbacks, as I will discuss further in Section 3.6

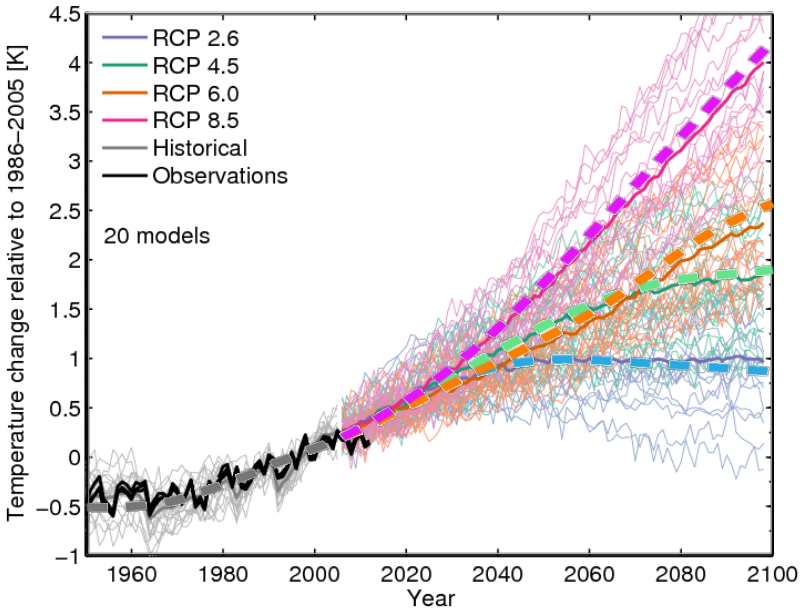
So in summary, the box-diffusion model is able to reproduce historical trends in CO_2 concentration and temperatures relatively well, with the exception of a period between the 1920s and 1940s which GCMs also have difficulty reproducing, possibly because this anomaly is due to natural variability (Meehl *et al.* 2004) or observation error Thompson *et al.* (2008). I also discussed how EBMs over-respond to volcanic forcing because of their lack of dynamics, as documented by Stone (2005). I will now do a different type of comparison with the box-diffusion model, this time against GCM output.

3.3 Comparison with GCMs

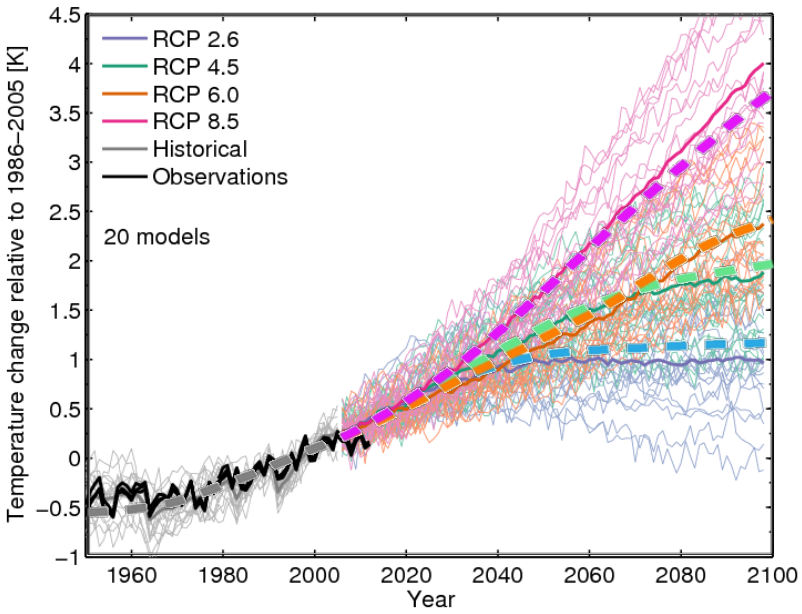
GCMs provide an additional source of data for simple models to benchmark themselves against. Reproducing GCM output with a simple model is easier than reproducing historical data, in part because some scenarios used to drive GCMs do not include volcanoes (Stocker 2011), which are not represented well by simple models.

In Figure 3.3 I compare the simple model against output from the latest round of the Coupled Model Intercomparison Project, namely CMIP5. In both panels in the figure, the box-diffusion model (dashed line) preforms relatively similarly to the multi-model means of the CMIP5 ensemble (thick solid lines), and well within the spread of the CMIP5 ensemble (thin solid lines). The box-diffusion model was tuned to the output of C^4MIP, so it is not that surprising that it tracks the multi-model mean of the ensemble well.

As expected, the model better tracks the output of CMIP5 when driven by concentrations in Figure 3.3a than when driven by emissions in Figure 3.3b. The CMIP5 ensemble presented here is a concentration-driven ensemble and so it is not surprising that our concentration-driven projections match them well (Taylor *et al.* 2012). The unofficial RCP emissions have been calculated using MAGICC, which has a different carbon cycle from us, and so the difference between the emissions-driven and the concentrations-driven projections arise simply from the difference between these two carbon cycles (Meinshausen *et al.* 2011c).



(a) Climate only comparison with CMIP5



(b) Coupled carbon cycle-climate box-diffusion model comparison with CMIP5

Figure 3.3: Comparing the box-diffusion model with CMIP5 model output. Thin lines are individual CMIP5 models (Taylor *et al.* 2012). Thick lines are the multi-model means across CMIP5. Dashed lines are box-diffusion model output. Colours correspond to the RCP emissions scenarios introduced in Section 2.5. Grey lines are historic model output. Black lines are historical observations (Lawrimore *et al.* 2011; Hansen *et al.* 2010; Jones *et al.* 2012). In Panel a, official RCP CO_2 concentration scenarios (Meinshausen *et al.* 2011c; van Vuuren *et al.* 2011) are used as inputs to the box-diffusion model and the carbon cycle is switched off. In Panel b unofficial RCP carbon emissions scenarios (Meinshausen *et al.* 2011c) are input into the box-diffusion model. In both panels, non-CO_2 forcings are added to the box-diffusion model as RFs. CMIP5 models use the official RCP concentrations as inputs (Taylor *et al.* 2012). Source: Adapted from an unpublished plot by Ed Hawkins (Met Office).

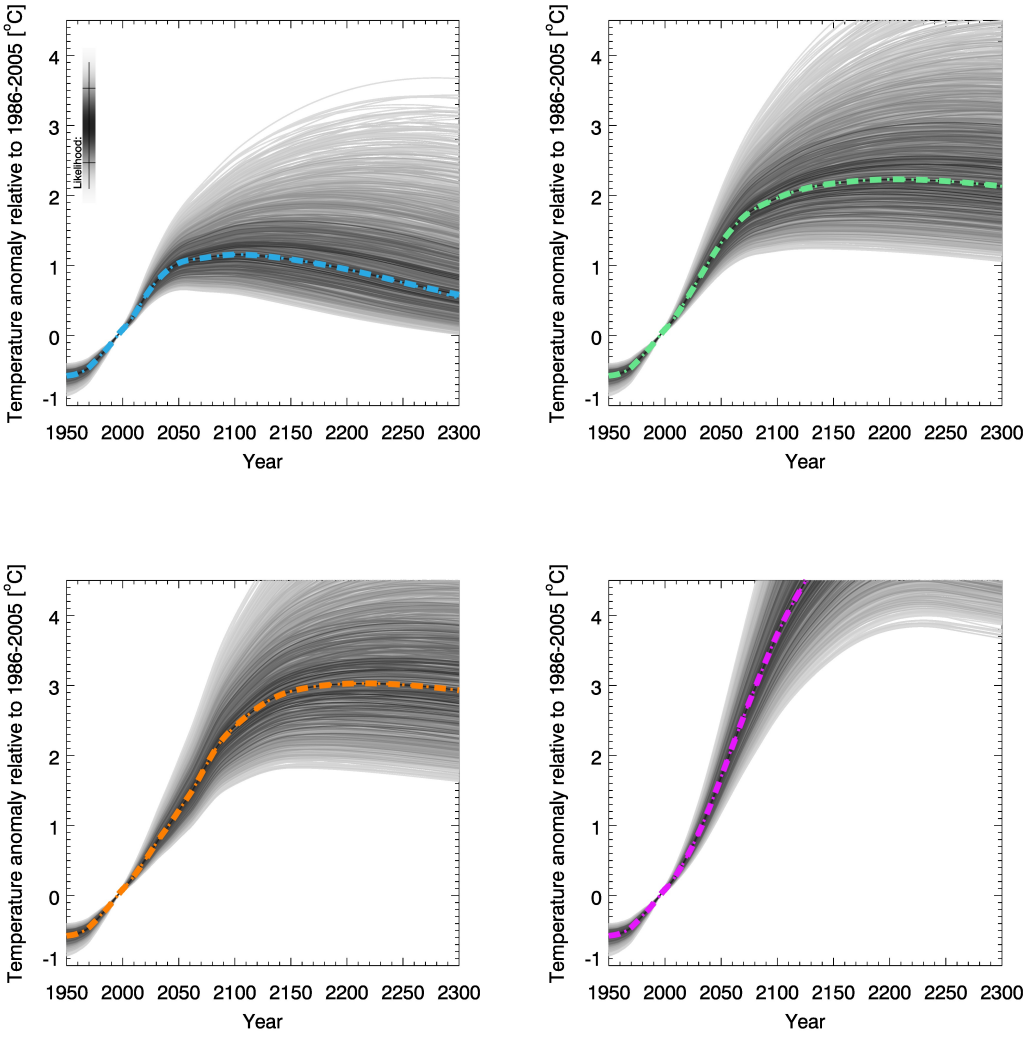


Figure 3.4: Global warming from the four RCP emissions scenarios (Meinshausen *et al.* 2011c) illustrated in Figures 3.3b and 2.8. The box-diffusion model is driven by CO_2 emissions non-CO_2 radiative forcings. The dotted lines are the best guess temperature responses from the box-diffusion model, as explained in Section 2.6, and at the colours correspond to those used in Figure 3.3. Thicker dashed lines use the asymmetric carbon cycle feedbacks scheme, while the thinner dotted lines use the symmetric carbon cycle scheme. The grey-shaded plume shows the 5-95% uncertainty on the dashed temperature response, as shown in the vertical scale bar. The vertical scale bar shows the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively.

The uncertainty in the box-diffusion model's projections in Figure 3.3b is shown in Figure 3.4, which uses the likelihood methodology described in Section 2.6.2.

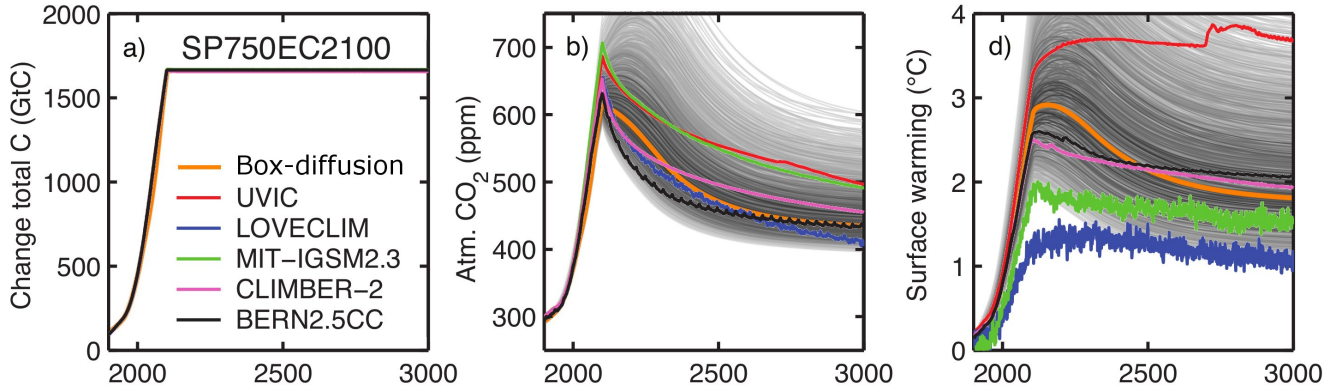
Given that the model tests well against GCMs in the case of rising or stabilising temperatures, I will now test it against EMICs and ESMs with their coupling between the carbon cycle and the climate, and I will also take this opportunity to test them in the case of falling temperatures.

3.4 Comparison with EMICs for a pulse emission

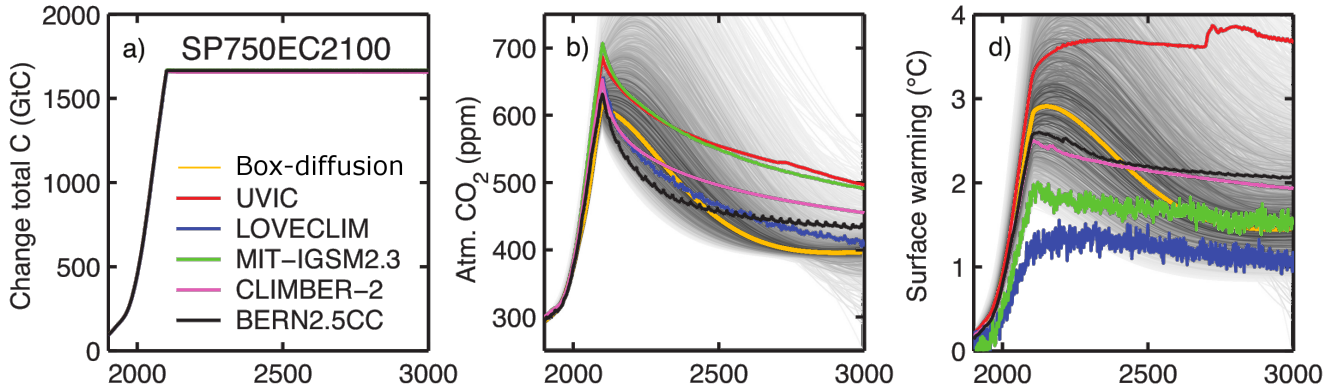
I have run the box-diffusion model driven by an approximation of the SP750EC2100 emissions scenario, which increases to 1700 GtC of cumulative emissions by year 2100, followed by a cessation of emissions (Plattner *et al.* 2008). In Figure 3.5 I have plotted this behind the same result from a number of EMICs in Meehl *et al.* (2007).

While the model response is roughly within the range of the EMICs, once emissions stop the carbon cycle removes CO_2 from the atmosphere faster than the EMICs, and consequently its temperature response also falls faster than the EMICs. This feature of the carbon cycle in the box-diffusion model is what makes me not consider its simulations as reliable once temperatures are falling. For this reason, throughout this thesis I have refrained from relying on output from the box-diffusion model once temperatures are falling in any of work I do in subsequent chapters in this thesis. I will discuss the limits of applicability of the model in more detail in Section 3.7. I will discuss the behaviour of the carbon cycle under falling global mean temperatures in more detail when discussing Figure 3.6.

Any discrepancy between the box-diffusion model and the EMICs before 2100 are likely due to the fact that I approximated the SP750EC2100 emissions scenario with the carbon emissions from RCP-4.5, which happens to have ~1700 GtC of cumulative emissions to year 2100. I altered RCP-4.5 to have zero emissions from year 2100 onwards, making it similar to SP750EC2100, though not identical until after year 2100.



(a) Asymmetric carbon cycle feedbacks in the box-diffusion model



(b) Symmetric carbon cycle feedbacks in the box-diffusion model

Figure 3.5: Comparing the box-diffusion model with EMICs. The upper panels use the box-diffusion model with symmetric carbon cycle feedbacks, and the lower panels use the box-diffusion model with asymmetric carbon cycle feedbacks. The left panels shows the emissions scenario that all models were forced with, but in units of cumulative emissions [GtC] rather than an emissions rate [GtC/yr]. The middle panels show the atmospheric CO_2 concentration in each model. The right panels show the global mean annual mean surface warming in each model relative to pre-industrial. The coloured lines correspond to the models indicated in the key. The yellow line is the box-diffusion model run with best-guess parameters. The grey shading is the uncertainty in the box diffusion model, as explained in Section 2.6. Source: Adapted from Figure 10.35 in Meehl *et al.* (2007).

In Figure 3.5 I considered the impact of 1700GtC of emissions; in Figure 3.6 I consider the impact of 1 ppmv-eq of emissions. I plotted the impact of a pulse of emissions in the box-diffusion model on (a) the CO_2 concentration anomaly; (b) the integrated radiative anomaly forcing divided by the instantaneous radiative forcing due to one ppmv of emissions; and (c) the global mean temperature anomaly. Over these plots, I copied the results from a number of EMICs and simple models from Joos *et al.* (2013).

In Figure 3.6a we see that the box-diffusion model has a different shaped response CO_2 concentration anomaly to the EMICs. The hump around 150 years in the response that we see in Figure 3.6a is caused by the carbon cycle feedbacks in the box-diffusion model.

The carbon cycle response in the box-diffusion model was designed with consultation with carbon cycle experts at the Hadley Centre in the years following the publication of Cox *et al.* (2000), which suggested that the carbon cycle feedback was large, and was not represented adequately in most simple carbon cycle models such as the Bern impulse response model (Joos *et al.* 1996). Since then, subsequent publications have suggested that the carbon cycle feedback is in fact smaller than that suggested by Cox *et al.* (2000) (Friedlingstein *et al.* 2006; Ciais & Sabine 2013).

As the carbon cycle feedback in the box-diffusion model is relatively large, the carbon sinks also take carbon out of the atmosphere relatively quickly in order to reproduce historical CO_2 concentration trends. This strong carbon cycle feedback is what causes this shape of the pulse response in Figure 3.6: first the strong sinks pull CO_2 out of the atmosphere quickly, and then the carbon cycle feedback starts emitting CO_2 as temperatures are increasing, and then once temperatures peak the carbon cycle feedback stop emitting CO_2 and so the strong carbon sinks take carbon out of the atmosphere rapidly bringing CO_2 concentrations back down again. At low likelihoods this phenomenon appears even more prominent, with low likelihood impulse response curves containing more carbon at 200 years than was originally injected.

Again, the box-diffusion model was not designed to respond well to pulse experiments

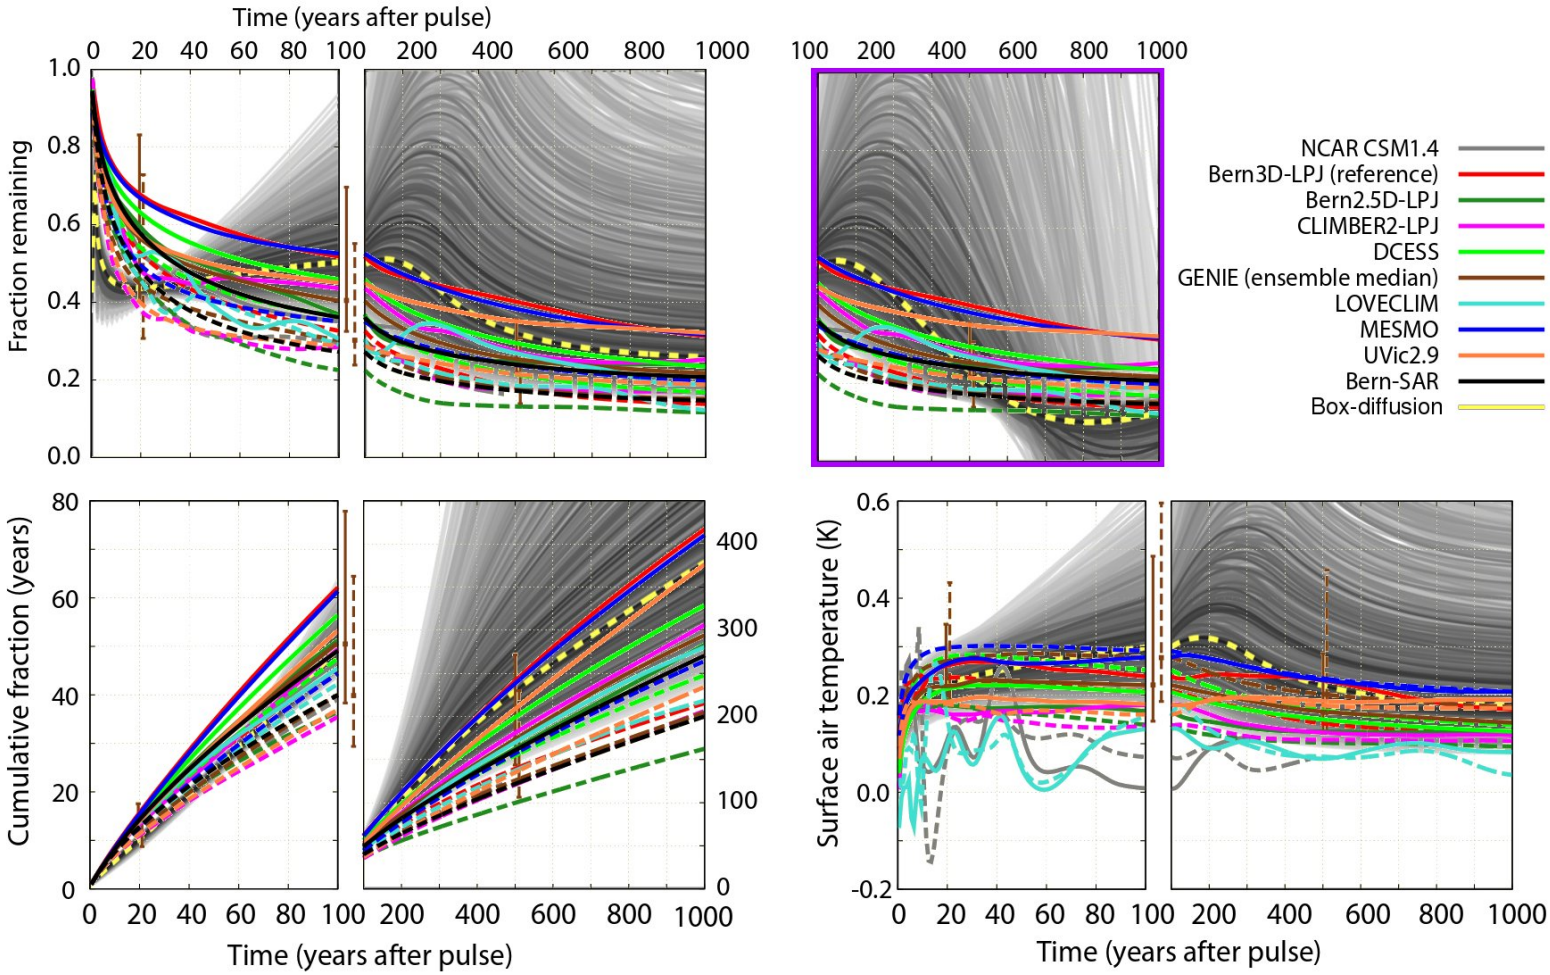


Figure 3.6: The evolution of the impulse response function for (a) CO_2, (b) integrated RF, and (c) temperature change, for an emission pulse of 100 GtC for a range of ESMs, EMICs, and simple models. For three models, Bern3D-LPJ (red), GENIE (brown) and MAGICC (green), an ensemble of simulations is available and the ensemble median and 5 to 95% confidence intervals are given by error bars for year 20, 100, and 500. The plot in the purple box uses the symmetric scheme for carbon cycle feedbacks in the box-diffusion, and is otherwise the same as the plot of the fraction of carbon remaining from 100 to 1000 years after the pulse shown directly to its left. The remainder of the plots use the asymmetric scheme for carbon cycle feedbacks in the box-diffusion model. Solid lines are for a pulse on a background concentration of 389 ppmv, and dashed lines are for a pulse on a pre-industrial background of 280 ppmv. I did not run the box-diffusion model for a pulse with a 389 background concentration as the model is not currently able to calculate the emissions required to maintain constant concentrations. The grey shaded plume represents the 5-95% confidence interval on the box-diffusion model as explained in Section 2.6. Source: Adapted from Joos *et al.* (2013).

followed by zero emissions. The model was tuned to historical concentration and temperature data (Allen *et al.* 2009c), as well as the C^4MIP ensemble under rising CO_2 concentrations (Friedlingstein *et al.* 2006), and so the box-diffusion model reproduces temperatures well when CO_2 concentrations are rising, but not necessarily when they are falling.

Figures 3.6b and c, are the same as Figures 3.8 and 3.9, but with different units and fewer models plotted. Thus, I will wait on discussing the remainder of Figure 3.6 until I discuss Figure 3.8 in Section 3.5.

As we do not have any historical data of the modern climate under rapidly falling CO_2 concentrations, it is difficult to know what the long-term response of the climate system will be once atmospheric concentrations start falling. As a result I am not that worried that the model doesn't fully reproduce the behaviour of other EMICs in this regime, particularly given that in this thesis I focus on peak temperatures so I do not use results from this regime by construction. Additionally, it is plausible that in the long term society will not take emissions fully to zero, but will instead produce some small quantity of GHG emissions from sources that are difficult or expensive to reduce emissions in such as CO_2 from air travel or nitrous oxide (N_2O) from fertiliser use, as I discuss in Section 5.1. This strategy could result in global means temperatures stabilising in the long term, if emissions are exactly matched with sinks.

The integrated radiative forcing shown in Figure 3.6, and the warming at a certain date, after a pulse of emissions, are both types of metric used to quantity the impact of different GHGs on global warming in different models. In the following section I will introduce metrics in more detail, before using them to compare the behaviour of the box-diffusion model against the behaviour of other models.

3.5 Metrics

Greenhouse gas metrics are well-defined quantitative measures that can be used to compare the behaviour of different GHGs between gases and between models. GHG metrics

measure the impact of emissions of GHGs on different climate change impacts. We can think of the cause-effect chain of climate change as running from emissions to damages via atmospheric concentrations, radiative forcing, climate change and climate impacts, as illustrated in Figure 3.7. Some metrics such as absolute global warming potential (AGWP) and global warming potential (GWP) compare the impact of GHGs emissions on radiative forcing (Harvey 1993), others such as global temperature change potential (GTP) (Shine *et al.* 2007) and integrated absolute global temperature change potential (iGTP) (Peters *et al.* 2011a) compare the impact of emissions on global mean temperatures, while metrics such as economic-damage index (EDI) (Hammitt *et al.* 1996), global damages potential (GDP) (Tol 1999) and global cost potential (GCP) (Johansson 2011) compare the impact of emissions on damages. I won't explain each of these metrics here, but will instead direct the reader to the citations above for a full explanation of iGTP, EDI, GDP and GCP. I will however introduce some of the metrics that are used to compare the impact of emissions on radiative forcing and global mean temperatures. The first of these metrics I will introduce is absolute global warming potential (AGWP) (Harvey 1993).

3.5.1 Absolute Global Warming Potential

AGWP is defined for a pulse emission of one kg of gas x as

$$\mathrm{AGWP}_x(\mathrm{TH}) = \int_{t=0}^{t=\mathrm{TH}} \mathrm{RF}_x(t)dt = A_x \int_{t=0}^{t=\mathrm{TH}} \mathrm{IRF}_x(t)dt \tag{3.1}$$

where x is the GHG, TH is the time horizon over which the AGWP is defined, $\mathrm{RF}_x(t)$ is the radiative forcing of gas x at time t (Harvey 1993).

We can break down $\mathrm{RF}_x(t) = A_x\mathrm{IRF}_x(t)$ where, $A_x = \mathrm{RF}_x(t = 0)$ is the radiative efficiency of gas x, where we define $t = 0$ as the time at which the pulse is emitted (Harvey 1993). $\mathrm{IRF}_x(t)$ is the fraction of the pulse emission of gas x remaining in the atmosphere at time t, or the impulse response function (IRF), or Greens function. Provided the pulse emission is small, and background conditions are approximately stable, A_x can be approximated as a constant (Harvey 1993).

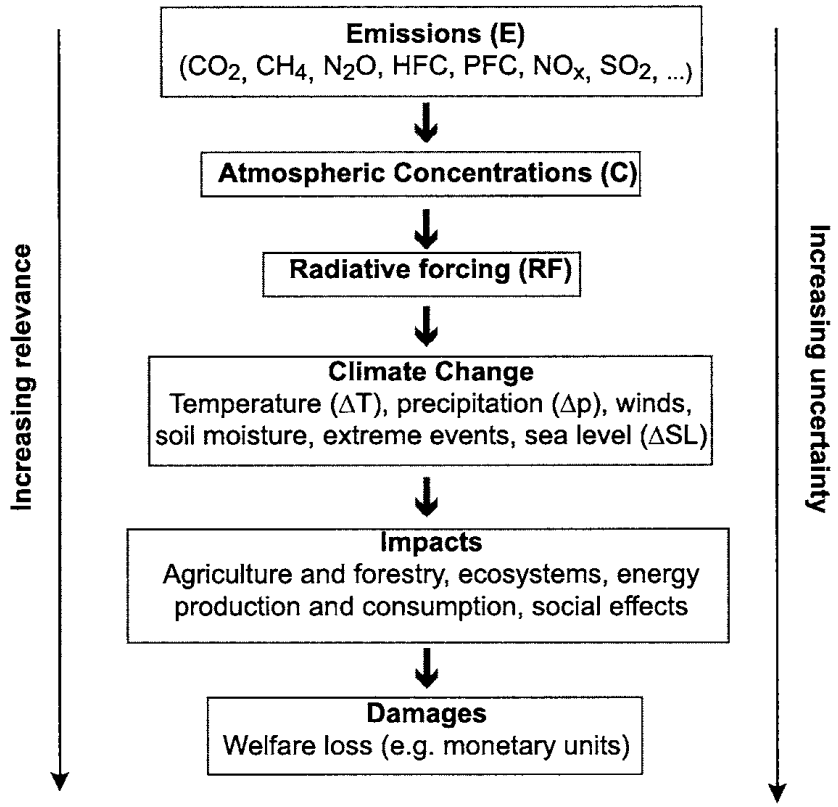


Figure 3.7: The climate change cause-effect chain, from emissions to climate change and damages. GHG metrics are a simple way of representing the propagation of effects down this chain from emissions to radiative forcing, global warming, climate change impacts or damages. Source: Fuglestvedt *et al.* (2003)

The choice of time horizon determines the limits of the integral of AGWP. The Kyoto Protocol, and in turn the European Emissions Trading Scheme (EU ETS), uses a 100-year time horizon, however this choice of time horizon is largely arbitrary and has little scientific basis beyond being the order of magnitude of the timescale to peak warming under a most scenarios in which temperatures peak at below a few degrees of warming, such as the RCPs shown in Figures 3.3 and 3.4 (O'Neill 2000; Shine 2009).

Unless otherwise stated, the CO_2 metrics used in this chapter are calculated relative to a background concentration of 378 ppmv following Joos *et al.* (2013). The box-diffusion model is not yet able to do the inverse calculation required to obtain the emissions required to stabilise CO_2 concentrations at 378 ppmv, so where I am attempting to calculate metrics relative to a 378 ppmv background I have done the carbon cycle component of the calculation relative to pre-industrial concentrations of 270 ppmv, with the carbon cycle feedback term turned on. I have then added 108 ppmv to the resulting CO_2 concentrations before using this new concentration time series to drive the climate component of the model, and calculating the metrics relative to the 378 ppmv background. The metrics derived for the remainder of the models in this section are done simply relative to a background concentration of 378 ppmv.

In Figure 3.8, I plot AGWP for CO_2 against the time horizon of the integral, both on logarithmic axes. I include the box-diffusion model, a number of simple models such as Bern-SAR, MAGICC6 and TOTEM2, a number of EMICs, and three ESMs which were not included in Section 3.4. I take the data for the models that are not the box-diffusion model from Table 4 in Joos *et al.* (2013). More information on each of the models and their setup can also be found in Joos *et al.* (2013). The figure shows that the best guess response from the box-diffusion model is mostly within the range indicated by the other models, including the ensembles shown in vertical coloured bars, at least beyond a 30 or 40-year time horizon. The primary difference between the box-diffusion model and the ensemble of other models is the gradient in the AGWP curve. The box-diffusion model has

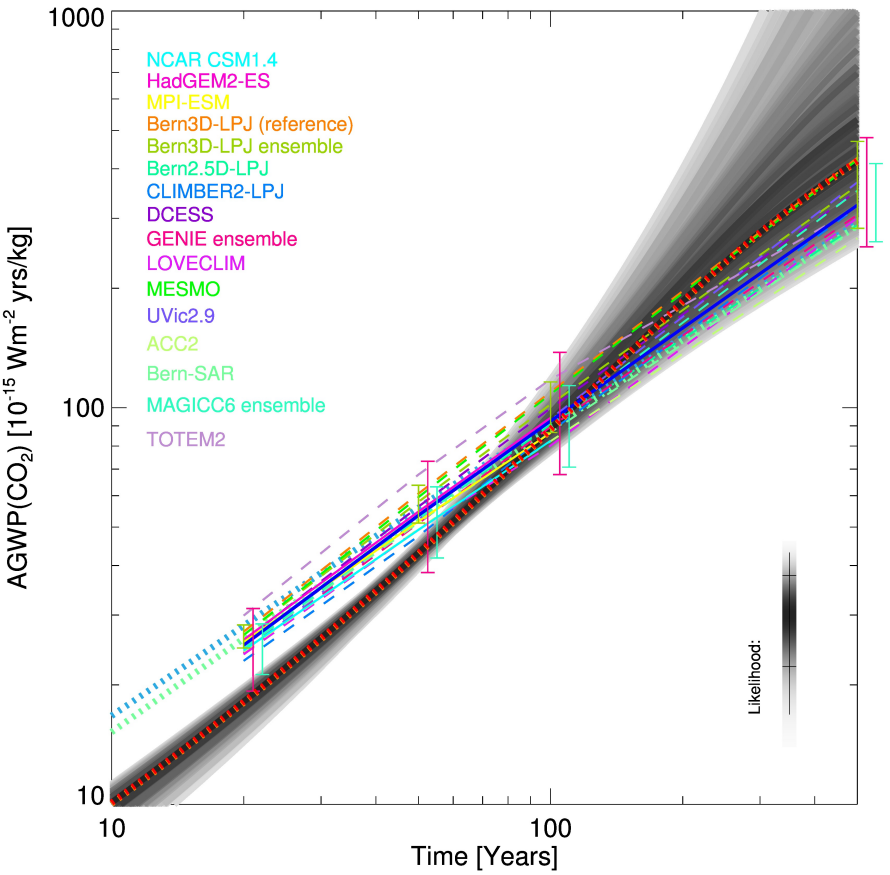


Figure 3.8: Absolute Global Warming Potential (AGWP) due to CO_2 in different simple models after a pulse injection of one ppmv-eq of CO_2. The blue dashed line is the result from Shine *et al.* (2005), and the green dashed line is the result from Fuglestvedt *et al.* (2010). The red dashed line is from the box-diffusion model. The orange dashed line uses the carbon cycle from the box-diffusion model and the linear conversion from CO_2 concentrations to forcing from Fuglestvedt *et al.* (2010). The solid lines are the results from a number of ESMs, EMICs, and simple models, and follow the colour key. The error bars are the ranges spanned by the ensembles of the model of the corresponding colour (the Bern3D-LPJ ensemble, the GENIE ensemble and the MAGICC6 ensemble). All data points in plotted in solid lines are taken from Joos *et al.* (2013). The solid blue line is the best estimate and the 5-95% confidence range on that estimate from Joos *et al.* (2013). The grey-shaded plume of dotted lines shows the 5-95% uncertainty on the red line, with the colours shown in the vertical scale bar in Panel b. The vertical scale bar in Panel b shows the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively.

a steeper gradient particularly between approximately the 40 and 400-year time horizon. This fits with the picture presented in Figure 3.6a, where the fraction of CO_2 remaining in the atmosphere after the pulse injection is seemingly too high compared with the remainder of the ensemble during this timeframe. Figures 3.6b and 3.8 are essentially an integral of Figure 3.6b given that we can linearise Equation 2.18 for small changes in C.

Another problem presented by Figures 3.6 and 3.8 is that the uncertainty range simulated by the box-diffusion model doesn't match the uncertainty range provided by the ensemble. This is because the parameters were selected to span the range of uncertainty in global warming under rising concentrations and temperatures. This is why the box-diffusion model is unable to span the range of uncertainty in the multi-model ensemble under declining concentrations and temperatures in Figures 3.6 and 3.8. If we wanted to adapt the box-diffusion model to behave like the impulse response function for the multi-model mean, and to span the range of uncertainty in the multi-model mean, then we would need to re-tune it with these as additional constraints, and produce a new set of likelihood parameters for the model.

A closely related metric to AGWP is GWP (Harvey 1993), which is defined for an arbitrary gas x as

$$\mathrm{GWP}_x(\mathrm{TH}) = \frac{\mathrm{AGWP}_x(\mathrm{TH})}{\mathrm{AGWP}_{\mathrm{CO_2}}(\mathrm{TH})} \tag{3.2}$$

This results in a dimensionless quantity that is useful for easily comparing the effects of different GHGs (Forster *et al.* 2007). The downside of GWP compared with AGWP is that GWP depends heavily on the AGWP of CO_2, which has been changing as we improve our understanding of the behaviour of CO_2, and in particular its impulse response function. When our best estimate of the AGWP of CO_2 changes, so does the GWPs of all the other gases (Joos *et al.* 2013).

I will discuss GWP in more detail in Chapter 6 when I introduce non-CO_2 GHGs into the box-diffusion model, but in this chapter I will continue the discussion of CO_2-only metrics by introducing absolute global temperature change potential (AGTP).

3.5.2 Absolute global temperature change potential

AGTP was developed from a desire to have a metric that was capturing impacts further down the cause-effect climate impacts chain (see Figure 3.7) than radiative forcing and its associated GWP and AGWP metrics (Shine *et al.* 2005). This is why the metric is defined as the impact of a pulse of emissions on the temperature after a given time horizon (Shine *et al.* 2005). Or put another way,

$$\mathrm{AGTP}_x(\mathrm{TH}) = \Delta T_x(t = \mathrm{TH}) \tag{3.3}$$

where AGTP_x is the global temperature potential due to gas x, $\Delta T_x(t)$ is the warming due to GHG x at time t after the initial emissions pulse at time $t = 0$. TH is the time horizon at which AGTP is defined. AGTP has units of °C kg^{-1}.

Like the relationship between AGWP and GWP, we can define a dimensionless global temperature change potential (GTP) for a given gas x as

$$\mathrm{GTP}_x(\mathrm{TH}) = \frac{\mathrm{AGTP}_x(\mathrm{TH})}{\mathrm{AGTP}_{\mathrm{CO_2}}(\mathrm{TH})}, \tag{3.4}$$

which has similar benefits relative to AGTP as GWP has relative to AGWP, as discussed in Section 3.5.1. In Chapter 8 I also use global peak temperature change potential (GPTP), which is GTP with the time horizon defined as the time of peak warming. I will discuss the GTP and AGTP of non-CO_2 GHGs in more detail in Chapter 6, and will focus here on the AGTP of CO_2.

In Figure 3.9 I plot $\mathrm{AGTP}_{\mathrm{CO_2}}$(TH) against TH. The first thing that I notice about the plot is that the lines do not start from the bottom-left of the plot and continue up to the top-right of the graph, as the did in the AGWP plot in Figure 3.8. This is because AGWP integrates the radiative forcing, while AGTP is the same as the temperature change induced by the pulse emission, without integration. iGTP integrates GTP over time, and so behaves more like AGWP, but I will not discuss iGTP further here.

While the box diffusion model (red dotted line) does not behave similarly it to the other simple models (coloured dotted lines), it does perform similarly to some of the other

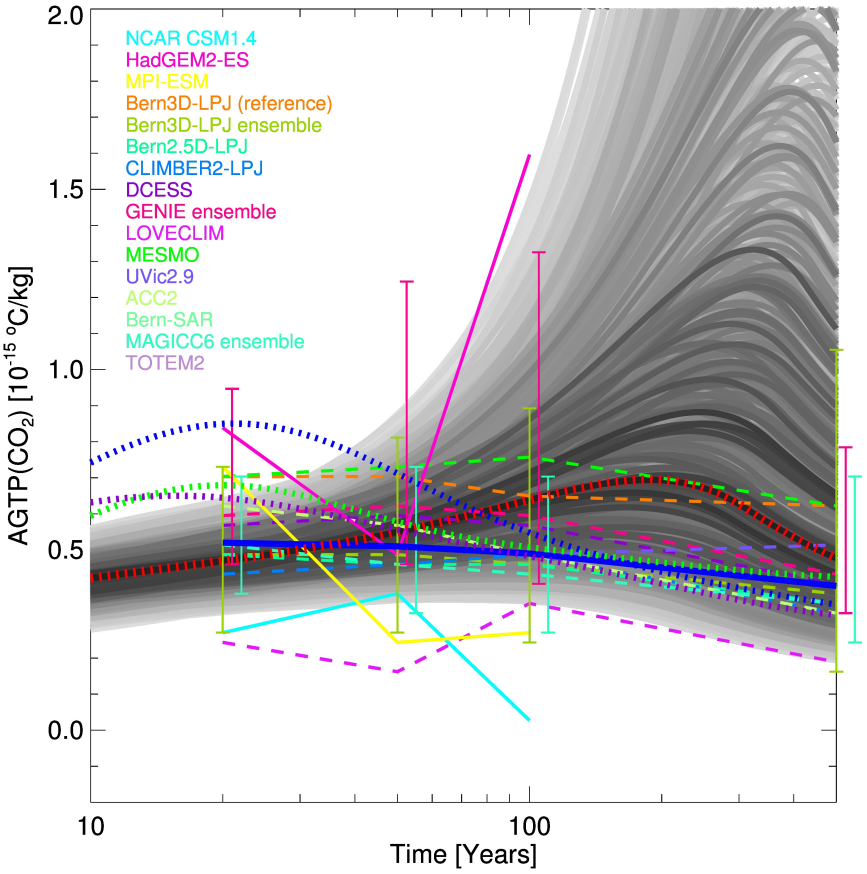


Figure 3.9: Absolute Global Temperature change Potential (AGTP) due to CO_2 in different simple models after a pulse injection of one ppmv-eq of CO_2. As in Figure 3.8, the blue dashed line is the result from Shine *et al.* (2005), and the green dashed line is the result from Fuglestvedt *et al.* (2010). The red dashed line is from the box-diffusion model. The purple dashed line represents a model that is a combination of the Bern carbon cycle model from Fuglestvedt *et al.* (2010) and the box-diffusion climate model. The solid lines are the results from a number of ESMs, EMICs, and simple models, and follow the colour key in Figure 3.8. The error bars are the ranges spanned by the ensembles of the model of the corresponding colour (the Bern3D-LPJ ensemble, the GENIE ensemble and the MAGICC6 ensemble). All data points in plotted in solid lines are taken from Joos *et al.* (2013). The solid blue line is the best estimate and the 5-95% confidence range on that estimate from Joos *et al.* (2013). The grey-shaded plume shows the 5-95% uncertainty on the red line, and uses the colour scale defined on the vertical scale bar in Figure 3.8.

models in the plot. The box-diffusion has a hump in the response around $\mathrm{TH} = 200$, which is visible in some of the other models such as LOVECLIM (in light purple near the bottom of the plot), HadGEM2-ES (in dark pink at the top of the plot), the warming models in the GENIE ensemble (in light pink bars), and the warmer models in the Bern3D-LPJ ensemble (in lime green bars). In the box-diffusion model, this hump in the AGTP response is largely driven by the hump in the fraction of carbon remaining in the atmosphere after the peak emission, as shown in Figure 3.6a.

The hump in AGTP in the four models listed above seems due in part to the behaviour of the carbon fraction remaining in Figure 3.6a, but also possibly due to internal variability. For example, the LOVECLIM model in Figure 3.6a has a clear hump in the response around 250 years, however it is unclear whether this is due to carbon cycle feedbacks or internal variability or a combination of the two. As another example, in Figure 3.9 the warming exhibited by HadGEM2-ES by year 100 is far larger than the remainder of the ensemble. A warming anomaly of this size seems likely to be due at least in part due to internal variability in the temperature response (Joos *et al.* 2013), however it looks also to be due in part to an increase in the atmospheric concentration of CO_2 between years 80 and 100, as illustrated in Figure 1a of Joos *et al.* (2013) (not shown here). These two effects are of course inter-related by carbon cycle-climate couplings. The spread in the response of most of the models in Figure 3.9 is likely to be due to a combination of their climate responses, and their carbon cycle responses, as shown in Figure 3.6 where the spread in the carbon fraction remaining alone is relatively large. Internal variability does not appear to play a key role in the response of most of the models that are not ESMs.

As in Figures 3.6 and 3.8, the uncertainty in the box-diffusion response seems likely to be too wide beyond the 100-year time horizon. This is due to the very large uncertainty in the CO_2 response shown in Figure 3.6. As discussed earlier, this is due to the fact that the model uncertainty was calculated for rising concentrations, and to get a more accurate spread in the uncertainty response here we would need to have tuned the model to an

impulse response scenario or similar. Again, there are no direct observations that I can tune to, and so I would have to tune to models, and hope that they are realistic.

The later hump and the wide spread in the uncertainty of the AGTP response in the box-diffusion is largely driven by its carbon cycle feedback response. Without the carbon cycle feedbacks enabled, the spread in uncertainty is narrower and the hump is earlier.

I would expect the wide spread in AGTP uncertainty from 100 years onwards to have some impact on TCRE estimates produced by the box-diffusion model, however not as much as we might initially expect from the plot above as the box-diffusion model does behave more like other models when forced using RCP-like scenarios as shown in Figure 3.3. This difference arises because the model was not tuned to an impulse-response scenario as discussed above. Indeed, when I compare the box-diffusion model's TCRE with the other estimates in the recent IPCC report in Chapter 8, I find that the box-diffusion estimate occupies the upper two-thirds of the range of the IPCC uncertainty range. Further discussion of this point and TCRE more broadly can be found in Section 8.2.

The box-diffusion AGTP CO_2 response is more within the envelop of other models than the box-diffusion model's AGWP response. This is likely to be because the AGWP plot only incorporates carbon cycle uncertainty (as well as any uncertainty in the temperature response that propagates into the carbon cycle via carbon cycle feedbacks), while the AGTP response incorporates both carbon cycle and climate uncertainty.

Figure 3.9 also shows the progression of simple impulse response climate models. The blue dotted line is from Shine *et al.* (2005), which peaks higher and falls faster than the remainder of the models. Since 2005, EMICs and ESMs have suggested that the temperature response to a pulse of radiative forcing is flatter than that of Shine *et al.* (2005), and the newer impulse response climate model proposed by Boucher & Reddy (2008) and used by Fuglestvedt *et al.* (2010) in the green dotted line with its two exponential functions better reproduces the response of the multi-model mean (Joos *et al.* 2013).

3.6 Symmetric and asymmetric carbon cycle feedbacks

Throughout most of this chapter I have presented the results for the carbon cycle in two different configurations. The asymmetric formulation, described by Equation 2.7 is what I have used primarily, but I have also shown the best-guess response, and in Figures 3.5 and 3.6 up to 2σ of uncertainty, on the response of the symmetric configuration of the carbon cycle feedback, described by Equation 2.16.

The easiest way to distinguish between the response of these two different formulations would be if we knew how fast CO_2 concentrations were to fall once emissions fell to zero. We do not have any observations for this future state, and so we must rely on other methods to compare between the two different responses. I can compare their responses in simulating past CO_2 concentrations, and I can compare them against simulations of future CO_2 concentrations in more complex models, as I have done in this chapter.

In Figure 3.2 I compare the response of the two different formulations in reproducing historical CO_2 concentrations. At first glance the asymmetric formulation appears worse at reproducing present day concentrations than the symmetric formulation, however the carbon cycle was tuned to the 20^{th} Century change in CO_2 concentrations using the symmetric configuration so if it were to be retuned using the asymmetric configuration the fit would be just as good. Note that I have not retuned the carbon cycle parameters throughout this thesis as the impact on warming and CO_2 concentrations is very small.

The other difference is between the late 19^{th} Century and the 1950s, when the asymmetric configuration reproduces historical concentration trends better than the symmetric configuration. This is likely to be due to the symmetric configuration taking large quantities of carbon out of the atmosphere immediately after major volcanic eruptions when temperature fall temporarily. Though there was initially discussion in the literature over whether volcanic eruptions cause CO_2 concentrations to rise or fall in the long term after a volcanic eruption (Jones & Cox 2001; Brovkin *et al.* 2010; Frölicher *et al.* 2011), a recent and more comprehensive study accounting for the many competing phenomenon affecting

CO_2 concentrations after historical volcanic eruptions has concluded that historic explosive volcanic eruptions are likely to have increased atmospheric CO_2 concentrations during the decade after the eruption, and so the response exhibited by the symmetric configuration appears to be unphysical (Frölicher *et al.* 2013). Because of this behavioural difference, the asymmetric configuration of the carbon cycle appears better at reproducing historical CO_2 concentration trends.

Comparing against GCMs simulating the RCPs in Figure 3.4, we see little difference between the two formulations, shown in thick dashed and thin dotted lines for the asymmetric and symmetric configurations. These simulations have only been run out until 2300, and running them later would probably show some differences when temperatures are falling faster. Our simulations of the RCPs seem robust to this change in the formulation of the carbon cycle feedback.

The difference between the two configurations is much more prominent in the impulse response simulations in Figures 3.5 and 3.6. The asymmetric formulation causes temperatures to fall slower and less far once they have peaked than the symmetric configuration. There is an even more prominent difference if we look in the tail of the likelihood profile at the end with large carbon cycle feedbacks. We see on the right hand side of the panel in Figure 3.6 with the purple box around fit that these low-likelihood models are falling in temperature very rapidly in this symmetric carbon cycle configuration. We also see this behaviour in the lower middle panel of Figure 3.5. Some of these models end up in runaway cooling, where the carbon cycle feedbacks keep removing carbon from the atmosphere because temperatures keep falling indefinitely driven by the logarithmic relationship between carbon concentrations and forcing in the model, with no regard to how fast the carbon cycle feedbacks can act nor where all of this carbon goes, which is clearly unphysical (Archer 2005). The asymmetric configuration does not display this behaviour in either Figure 3.5 or 3.6. In both plots the asymmetric configuration better matches the behaviour of the other models as well.

Both the comparison against historic concentrations, and the impulse response comparison against other models suggests that the asymmetric configuration is a more accurate representation of carbon cycle feedbacks than the asymmetric configuration. Thus I will use this configuration in the rest of this thesis.

In summary, despite not being designed to function well under impulse response scenarios, the box-diffusion model with the asymmetric carbon cycle configuration stayed mostly within the range of behaviour of other EMICs and ESMs, though it behaved differently in a number of ways which I will use to illustrate the limits and applicability of the model in the following section.

3.7 Limits and applicability of the model

"All models are wrong, but some are useful," wrote George Box (1976). This applies to any model, and the box-diffusion model is no exception. In this section I discuss some of the ways in which the box-diffusion model is wrong, and then go on to describe some of the ways in which the model can be useful.

3.7.1 Limits of the box-diffusion model

In this section I will discuss some of the regimes in which the model is unreliable.

Falling temperatures One of the key limits of the box-diffusion model is that it becomes less reliable shortly after global mean temperatures peak. As I have mentioned before, the box-diffusion model was designed and tuned to scenarios with rising concentrations using the output from the C^{4}MIP ensemble. Thus, the output from the box-diffusion model is unreliable when global mean temperatures are falling.

High temperatures and non-linearities The coupled climate-carbon cycle system has a number of potential "tipping elements" that could make the system's behaviour non-linear in certain regimes (Lenton *et al.* 2008). Simple climate models do not generally represent

these non-linearities. As a result we cannot trust the results of simple climate models in these non-linear regimes, though a simple model will predict behaviour relatively well up until these non-linear regimes.

Which regimes will behave nonlinearly is currently beyond our level of scientific understanding. The risk of non-linearities is thought to increase the further the climate system is taken from equilibrium (Schneider 2004). As the box-diffusion model does not represent these non-linearities it is likely to become an even poorer model of reality the further it is taken from equilibrium. Unfortunately this may be an issue that affects all of climate science as it is unclear whether these non-linearities are well-represented in GCMs, though they are generally better represented in GCMs than in simple models where they are often not included at all. As a result the box-diffusion model would not be able to reproduce an abrupt change caused by a non-linearity such a shutdown of the MOC, retreat of the arctic sea ice, a collapse of the Greenland or West Antarctic ice sheets, a rapid and irreversible change in vegetation cover, or an abrupt change in the atmospheric and ocean-atmosphere regimes (Meehl *et al.* 2007).

Because of the model's inability to model the potential for abrupt or non-linear climate change, I would trust the output of the model less for higher temperatures. Up to 2°C, it is unlikely that we will find large non-linearities, however this risk increases to the point that I would not trust the output of the model above the 3–4°C range because of the non-linearities not represented in the simple model (Collins *et al.* 2013).

Inter-annual variability The box-diffusion model is a simple model that simulates trends in global mean annual mean climatic temperatures under forcing from GHGs. By annual climatic temperatures, I mean the 30-year running mean of global mean temperature (Baede 2007). The box-diffusion model is unable to reproduce any variability because of its lack of the relevant mechanisms such as dynamics. As such, the output from the box-diffusion model is unable to simulate the variability of the climate, or predict any changes in that variability.

Because of the lack of inter annual variability in the model, when I mention 'peak warming' or the 'maximum rate of warming' I don't mean the actual peak temperature or the actual maximum rate, I mean the climatic average maximum, or the maximum in the 30-year average.

Climate processes and regional responses Simple models such as the box-diffusion do not include almost every single one of the many phenomena that combine to produce the climate in a particular country or region. For example, the box-diffusion model does not explicitly model phenomena such as atmospheric and ocean circulations, the equator-poles temperature gradient, Rossby waves, ocean gyres, cloud formation, atmospheric blocking, land mass, orography, precipitation, the cryosphere, and many other phenomenon. As a result the box-diffusion model is unable to reproduce any of the phenomenon that govern regional climate change such as convergence zones, oscillations such as the Madden Julian oscillation, the quasi-biannual oscillation, Atlantic ocean modes, El Nino and related teleconnections, annual and dipolar modes, monsoons, and many other phenomenon. The box-diffusion model is so simple that it is unable to make projections on national, regional or sub-global scales, and even on a global scale it is only able to project CO_2 concentrations and temperatures. Thus we can use such a simple model if, and only if, we are concerned with globally-averaged phenomena. While these sub-global scale climate phenomena are important for determining the Earth's climate, for the most part I do not need to explicitly include them in my simple climate model because their affect on global average climate has been incorporated by tuning the box-diffusion model to the results of more complex models.

If the climate moves into a non-linear regime these phenomena may become far more important for determining the global average climate, however as discussed above, our model is not expected to represent reality at all well the more non-linear the climate regime.

Non-temperature climate impacts Climate change takes the form of changes in a number of variables with respect to time. Temperature is just one of these variables, but climate change is also predicted to impact upon rainfall, wind speeds, the cryosphere, extreme weather, and atmospheric and oceanic circulations, among other changes (Hartmann *et al.* 2013; Rhein & Rintoul 2013; Comiso & Vaughan 2013). As the climate changes, many of these changing phenomenon are thought to have socio-economic impacts through droughts, heat waves, wildfires, species migration and extinction, changing crop yields and changing disease risk areas due to malaria and dengue fever, among other impacts (Fankhauser *et al.* 2001). Because the box-diffusion model contains only a simple energy balance model (EBM), it is only able to simulate changes in global mean temperatures, and cannot make statements about these other impacts, despite the fact that they are what will ultimately cause much of the damages due to climate change. Rising temperatures are simply a trigger for these more damaging impacts.

Emulating GCMs As I mentioned at the beginning of this section, "All models are wrong", and GCMs are no exception. For example the CMIP5 ensemble contains "significant errors in the model simulation of clouds" (Flato & Marotzke 2013; Boucher & Randall 2013); sea ice models have mostly failed to reproduce the observed trend in annual minimum Arctic sea ice extent (Stroeve *et al.* 2012; Winton 2011); and contemporary carbon cycle models tend to underestimate the land carbon sink (Eby *et al.* 2013). Because the box-diffusion model is essentially simply an emulator of GCMs, it is possible that the box-diffusion model is under- or over-estimating the extent of global warming because of not aggregating in these climate processes that are poorly represented in GCMs.

Probabilities The likelihood methodology used to quantify uncertainty in the box-diffusion model takes a frequentist approach, and is related but different to the probabilistic approach more commonly applied in climate modelling (Pawitan 2001).

Because probability distributions depend heavily on the choice of prior (Frame *et al.* 2005), the box-diffusion model was designed to use a more robust and objective indicator of uncertainty, namely likelihood profiling. Likelihoods say nothing about the chance of an event occurring in the future, but simply say how well a possible future event would fit with historical observations in a given model. As such the likelihoods used throughout this thesis do not include model or structural uncertainty, nor initial condition uncertainty. They also do not require us to incorporate our prior beliefs into uncertainty distribution, apart from indirectly through the structure of the model (Pawitan 2001). The advantages and drawbacks of likelihoods are discussed in more detail in Section 2.6.

Crucially, unlikely probability distribution functions, the method of quantification of uncertainty used here makes no statement about the chances of a certain climatic projection occurring in the future. Instead we use likelihood profiling which simply states how well a potential future projection fits with the constraints discussed in Section 2.6.2 given the box-diffusion model.

3.7.2 Behaviour of the box-diffusion model relative to other simple models

Above I have discussed some of the limits of the box-diffusion model. In this section I will discuss traits of the box-diffusion model that are potentially within the uncertainty bounds of the ensemble of GCMs, but that behave differently from other simple climate models.

A key way in which the box-diffusion model differs from other simple climate models is that it has an explicit carbon cycle feedback term, as shown in Equation 2.7. This means that it has a strong carbon cycle feedback response relative to most other simple models (Joos *et al.* 2013). The box-diffusion model's carbon cycle feedback term was tuned to the output of the C^4MIP ensemble, and so it largely reproduces the magnitude of the carbon cycle feedback terms of the models there.

Because of the additional carbon emitted from the carbon cycle feedbacks, In order to reproduce historical atmospheric CO_2 concentration trends the carbon sinks in the box-

diffusion model are also stronger than those in conventional simple models. This means that once emissions fall to zero and temperatures stabilise, CO_2 concentrations the model will fall more rapidly than in more complex EMICs and ESMs, as already discussed in Section 3.7.1.

We also see that the fraction of carbon remaining in the atmosphere after a pulse emission is different in the box-diffusion model than it is in the other models in Figure 3.6. The box-diffusion model response falls faster in the first decade, and then is flatter and even increases unlike the other models over the next century, before falling similarly to the other models for the next millennia but starting from a higher atmospheric fraction and so remaining higher than the other models. At the warmer tail of the likelihood profile, the behaviour of the models appears to be more unrealistic, with some of the carbon cycle feedbacks in the models being so strong that the fraction of carbon remaining in the atmosphere is greater than unity. The nearest occurrence to this phenomenon that I have observed in the literature is the close to flat CO_2 concentration response to zero emissions that we can observe in the GCMs with stronger carbon cycle feedbacks, for example in HadCM3LC when preceded by the Special Report on Emissions Scenarios (SRES) A2 scenario until 2100 (Lowe *et al.* 2009), however even here concentrations fall by more than 50 ppmv during the 22^{nd} Century.

In order to constrain the impulse response we would need an additional constraint, such as the evolution of the fraction of atmospheric C^{14} after the nuclear test ban treaty came into effect in 1963 (Hua & Barbetti 2004), in combination with a constraint on the quantity of C^{14} that had been cycled out of the atmosphere through gas exchange, for example between the atmosphere and ocean.

This discrepancy in fraction of carbon remaining in the atmosphere after a pulse emission is potentially the largest problem with the box-diffusion model, and if there are problems with any of the results presented in this thesis I would imagine that that this could be the most likely place that they might be coming from.

Nonetheless, despite this discrepancy in the impulse response, the box-diffusion model behaves very similarly to the multi-model mean of the CMIP5 models in Figure 3.3, and then has a relatively flat temperature response for the next couple centuries after 2100, in agreement with the results from EMICs and ESMs suggesting a relatively flat post-peak temperature response (Matthews & Caldeira 2008; Solomon *et al.* 2009; Gillett *et al.* 2011). The RCPs were designed to span the range of possible emissions scenarios over the coming century, from extreme mitigation to business as usual, and so because the box-diffusion model behaves so similarly to the CMIP5 multi-model mean across as wide a range of scenarios, I am inclined to assume the difference in impulse response doesn't have a major impact on the ability of the box-diffusion model to simulate the response of more complex models under the range of relatively plausible future emissions scenarios.

Given the various limitations and drawbacks of the model, and the advantages already outlined in Section 1.5, what is the box-diffusion model useful for, and how should it be used?

3.7.3 Applicability of the box-diffusion model

The box-diffusion model is a simple tool that will never be able to match more complex GCMs in terms of its ability to represent components of the climate system accurately. Nonetheless, it is still a very useful research tool.

The box-diffusion model, like other simple models, is computationally cheap and very flexible, which means it can be used to run simulations of future climates far faster than can be done using more complex models. As such, it is an invaluable initial tool for use before a GCM in some areas of research. Because of its simplicity it can also provide insights into processes and phenomenon that might have been previously overlooked (Stocker 2011). As such, the box-diffusion model can be useful valuably as an initial tool to highlight areas in which it would be useful to run GCM simulations.

Additionally, the box-diffusion model can be used to explore uncertainty. The box-diffusion model can easily provide likelihood profiles quantifying parameter uncertainty

in future projections, an output that can be difficult or impossible to obtain from more complex models without significant computational cost (Stocker 2011). Because the box-diffusion model uses constraints directly to quantify uncertainty as discussed in Section 2.6.1, these likelihood profiles have the potential to be more objective than quantifications of uncertainty based on expert elicitation or using subjective priors.

It is in the uses described in this section that I will employ the box-diffusion model over the remainder of this thesis. I will begin by exploring emissions targets and their impact, assuming they are successfully achieved, on global mean temperatures in the next chapter.

Chapter 4

Emissions targets

In this chapter I present some of the results of my thesis, which have been obtained using the box-diffusion model I presented and tested in Chapters 2 and 3. I present a comparison of six different types of emission target, including cumulative targets and annual targets. I compare the impact of these targets on peak warming and the rate of warming. I published some of these findings, and those of Chapter 5 in Bowerman *et al.* (2011).

4.1 Motivation

This work was carried out around the time of the Copenhagen climate change conference. Heads of state from around the world were meeting to discuss emissions targets. There was much discussion of annual targets, such as 2020 targets and 2050 targets, but very little discussion of other types of emission target such as cumulative targets (UNFCCC 2009). Several NGOs were discussing the date at which emissions peak, but again with little to no emphasis on cumulative emissions (Huntingford *et al.* 2012). In the Copenhagen Accord, leaders agreed to limit global warming to 2°C, so in Section 4.2, I investigate which of these proposed target types is most relevant to the goal of keeping global warming below this 2°C target.

Article 2 of the UNFCCC states that the mitigation of climate change *"should be achieved within a time-frame sufficient to allow ecosystems to adapt naturally to climate change."* (United Nations 1992) Several authors have hypothesised that the ability of some

ecosystems to adapt will depend not only on the magnitude of warming but also the rate of warming (Root *et al.* 2003; Leemans & Eickhout 2004; Jump & Penuelas 2005). In Section 4.5 I attempt to determine which of these types of target is most relevant to the rate of warming.

4.2 Comparing emissions metrics

In this chapter and the next, I will be addressing only the issue of CO_2-induced warming. This is a central but not exhaustive component of dangerous anthropogenic interference with the climate system. Most multi-gas pathways of future radiative forcing currently in the literature describe a total anthropogenic warming that either approximately equals or exceeds CO_2-induced warming (Nakicenovic *et al.* 2000; Meinshausen *et al.* 2011c). This is due to the warming effect of non-CO_2 greenhouse gases usually equalling or exceeding the cooling effect of aerosols. Hence, avoiding dangerous levels of CO_2-induced warming is a necessary, albeit not always sufficient, condition for avoiding potentially dangerous anthropogenic interference in the climate system.

Most of the largest non-CO_2 anthropogenic forcing agents are distinct from CO_2 in having much shorter effective lifetimes in the climate system (Smith *et al.* 2012). The exception is nitrous oxide, which has an atmospheric lifetime comparable to, albeit different from, CO_2 and, crucially, longer than the response-time of the physical climate system (Smith *et al.* 2012). Hence it could be argued that nitrous oxide emissions should be considered in the same framework as CO_2, and throughout this chapter and the next, CO_2 could be replaced by the combination of CO_2 and N_2O. Predicting long-term emissions of N_2O, given current rapid developments in agricultural technology, is even more difficult than predicting CO_2 emissions, so for the sake of simplicity we focus on CO_2 alone, and do not consider the issue of what fraction of these long-term emissions might be made up of N_2O, although we note that in other papers considering the impact of emission over the very long term, this fraction is substantial (Anderson 2008).

I begin by considering the impact that different global CO_2 emissions target would have on peak warming if they were achieved. I only consider CO_2-induced warming here, and consider non-CO_2 GHGs from Chapter 6 onwards.

I compare the performance of a range of emissions and cumulative emissions metrics in estimating peak CO_2-induced warming. I do this comparison by constructing an initial set of 399 different emissions pathways, 395 of which have been randomly selected from the possible pathways outlined in Section 2.5.1. The remaining four are the CO_2 pathways from the RCPs outlined in Section 2.5.1. Once I have randomly selected the 395 emissions pathways, I use the box-diffusion model described in Chapter 2 to estimate quantities such as the most likely peak warming for each pathway. I use these results to analyse the usefulness of each of six emissions metrics of interest. I consider cumulative carbon emissions a) from pre-industrial times to the time of peak warming and b) from year 2010 to year 2050. I also consider the actual emissions rates at c) year 2020 and d) year 2050. Additionally I consider e) the peak emissions rate and f) the year in which emissions peak.

The performance of each emissions metric is shown in Figure 4.1, where the emissions metrics are plotted against the peak warming. As introduced in Section 2.6.2, I have used high response parameters in red and orange, best-guess response parameters in black and grey, and low response parameters in light and dark blue.

Based on the metrics presented in Figure 4.1, I conclude that for cases with no emissions floor, the strongest correlation across all pathways occurs between peak warming and the cumulative emissions from pre-industrial times to the time of that peak warming, as shown in Figure 4.1a. The correlation is almost as strong if cumulative emissions out to 2500 are considered (shown in black symbols in Figure 5.2a) because the vast majority of the emissions in these zero emissions floor pathways have occurred by the time of peak warming.

In Figure 4.1a, the correlation for each of the response parameters is very high, apart from the crosses representing RCP-3PD. RCP-3PD has net negative emissions from the late

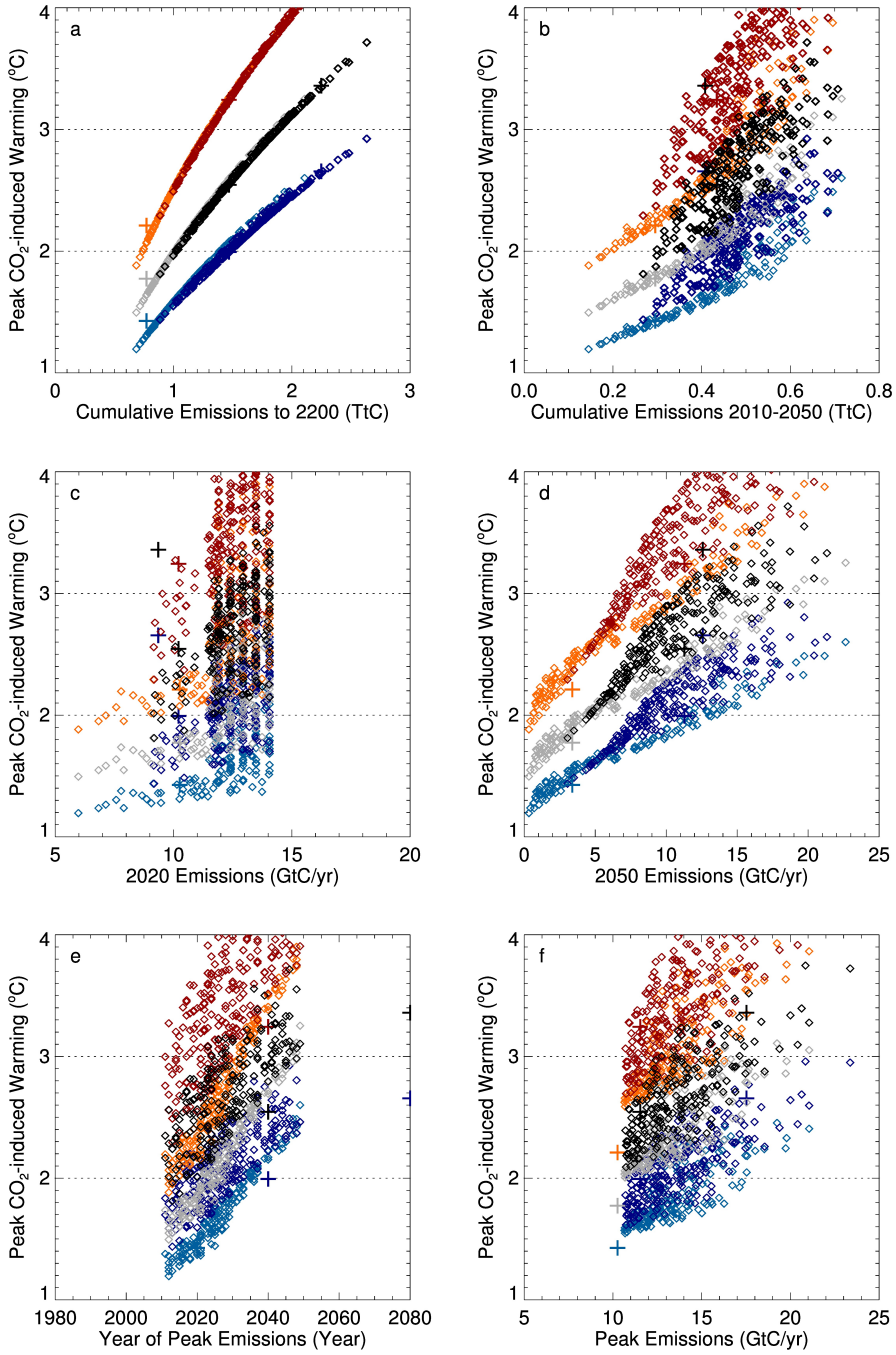


Figure 4.1: Scatter plots showing the relationship between peak CO_2-induced warming and various global carbon emission metrics for 399 emission pathways. The x-axis for each panel shows: (a) emissions time-integrated between 1750 and 2500, (b) emissions time-integrated between 2010 and 2050, (c) emissions in the year 2020, (d) emissions in the year 2050, (e) the year in which emissions peak, and (f) the peak or maximum in emissions. Darker diamonds (in black, dark red, and dark blue) indicate emission pathways in which the maximum rate of emissions decline is less than 4%/yr; lighter diamonds (in grey, orange and light blue) indicate the converse. The red and orange symbols are from models at the warmer upper bound of the 17%-83% confidence interval in peak warming. Blue symbols are from models at the cooler lower bound of the 17%-83% confidence interval in peak warming. Grey and black symbols are from models with the best guess peak warming, as discussed in Section 2.6. Crosses represent the RCPs.

21st Century onwards, which moves the crosses slightly to the left of where they would be if we only considered cumulative emissions to peak. This issue is discussed further in Section 4.3.

For Figures 4.1 and 4.4, grey, orange and light blue diamonds, represent emission pathways with a peak rate of decline of at least 4% per year and not more than 10% per year. Emissions pathways with rates of decline greater than 10% per year are not included here. These rapidly declining pathways generally appear to the right and below the black, red and dark blue crosses, which represent emissions pathways with peak rates of decline between 0% and 4%. For a given peak warming, and so for a given cumulative total, smooth pathways with a faster rate of emissions decline will have relatively more of their cumulative total emitted sooner than for pathways with slower rates of decline. As a result these rapidly declining pathways will have higher cumulative emissions between 2010 and 2050, higher 2020 emissions, higher peak emissions, and a later year of peak emissions. This effect also holds for 2050 emissions above 5GtC/yr in these smooth pathways. For smooth emissions pathways with cumulative emissions less than 1TtC, corresponding to a peak warming of 2°C, 2050 emissions occur once the pathway has been declining exponentially for a considerable period, and so rapidly declining pathways will have relatively lower emissions in 2050.

The correlation between peak warming and cumulative emissions between year 2010 and year 2050, which is the emissions metric used by The German Advisory Council on Global Change (WBGU 2009), is plotted in Figure 4.1b. The correlation here is not as good as in Figure 4.1a. Emissions before 2010 are not allowed to vary across emissions pathways, so there can be no contribution to the spread in peak warming from this historical time period. The majority of the spread comes from the variation in post-2050 emissions, which will have a significant impact on peak temperatures, but which by definition are not included in 2010-2050 metrics. I see that in cases where post-2050 emissions are small, the spread is much tighter, as shown by those pathways with cumulative emissions less

than 0.3TtC between 2010 and 2050. This is because the majority of future cumulative emissions in these pathways are emitted before 2050. This means that up to approximately 0.3 TtC, the cumulative emissions between 2010-2050 has some skill in predicting peak CO_2 induced warming, but this skill is reduced for higher cumulative targets.

I also consider whether there is predictive skill in using the actual emissions rates at a particular year; here I use year 2020 and year 2050. The former year is chosen because most of the Copenhagen Accord emissions reduction pledges are quoted for this year (UNFCCC 2009). The latter is chosen because several reduction targets for 2050 have also been presented (G8 2009; Rogelj *et al.* 2010, 2011a). Figure 4.1c and Figure 4.1d show 2020 and 2050 emissions against peak warming. As we confirmed in Huntingford *et al.* (2012), I found that emissions in the year 2020 are not a good indicator of peak warming, presumably because they are largely a function of current emissions, and are not a key determinant of cumulative emissions.

For Figure 4.1d, year 2050 emissions do seem to be a relatively good indicator at lower rates of emission, particularly at values that cause approximately two degrees of warming or less. However I found 2050 emissions to be a less good indicator of peak resultant warming at higher rates of emission. This is in part a consequence of the choice of pathways, as I have considered only smooth pathways with a single maximum and exponential tails (see 2.5.1 for more detail on how the emissions pathways are chosen). A wider range of functional forms to describe emissions pathways would be expected to reduce the strength of the relationship between 2050 emissions and peak warming.

I also compare the peak emissions rate and the year of peak emissions with the peak CO_2-induced warming. As shown in Figure 4.1e and Figure 4.1f, the spread is very large for both of these metrics, and there is little correlation, save the rapidly declining emissions pathways appearing to the right of the slowly declining pathways, as explained earlier.

4.3 Extreme emissions pathways

In order to test the concept of cumulative emissions further, I have developed a set of extreme and unphysical emissions pathways to test the concept.

The extreme emissions pathways that I have created follow historical emissions up until year 2000, and have zero emissions after year 2200, but have have a cubic interpolation between randomly selected emissions rates each decade between year 2000 and 2200 (Park & Schowengerdt 1983). The randomly selected emission rates, $E_{i,j}$ are selected for emissions pathway i, in year $2000 < 2000 + 10j < 2200$ using the following equation

$$E_{i,j} = \left(\xi_{i,j}\left(E_{\max} - E_{\min}\right) + E_{\min}\right)\Xi_i, \tag{4.1}$$

where $0 \leq \xi_{i,j} \leq 1$ is a random number chosen for each decade and emissions pathway; $E_{\max} = 100$ GtC; $E_{\min} = -50$ GtC; and $0.1 \leq \Xi_i \leq 1$ is a random number chosen for each emissions pathway. This equation can be considered as a bounded random emissions profile generator, multiplied by a random number to vary the total cumulative emissions in each profile. The yellow and red curves in Figure 2.8 are examples of these extreme emissions pathways. Having described the extreme emissions pathways, I will now describe how I used them to test the cumulative emissions concept.

In Figure 4.2, I have reproduced Figure 2 of Allen *et al.* (2009c) but using a subset of the smooth emissions pathways described by Equation 2.25 in blue, and including some extreme emissions scenarios in red. The relationship between peak warming and the cumulative emissions integrated from 1750 to the time of peak warming remain the same in the different types of pathways. I have also included the RCPs in the figure in triangles, which also lie on the curve.

The only two pathways that appear as anomalies in Figure 4.2 appear about a quarter of a °C above the curve containing the remainder of the pathways. These pathways are unique from the other pathways in that they have heavily net negative cumulative emissions by the 22$^{\text{nd}}$ Century when integrated from 1750. They then have very large emissions (and

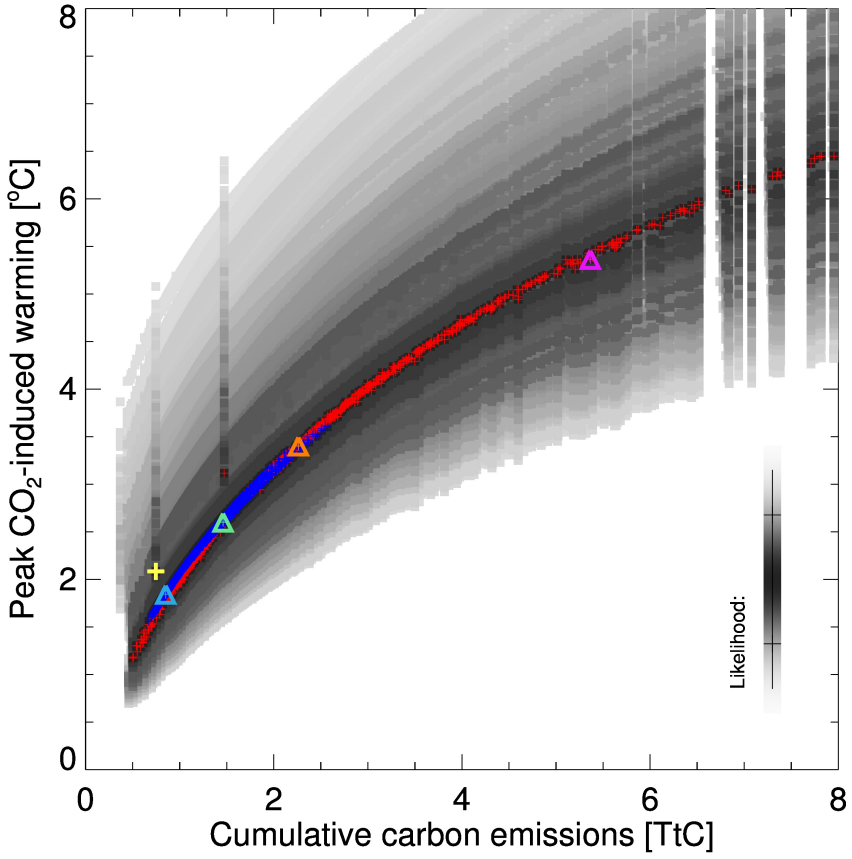


Figure 4.2: Peak CO_2-induced warming for different emission pathways as a function of total cumulative emissions integrated from 1750 to the time of peak warming as simulated by the box-diffusion model. Figure 2.8 shows a sample of the emissions pathways used to create this plot. Coloured symbols correspond to best-fit values of simple climate model parameters, with each cross corresponding to a single scenario. Coloured triangles correspond to the RCPs as shown in Figures 3.4 and 2.8. Blue crosses correspond to the idealised emissions scenarios described by Equation 2.25 and shown in grey and blue in Figure 2.8. Red crosses correspond to the extreme emissions scenarios described in Section 4.3. One of the red scenarios is plotted in Figure 2.8. The yellow cross corresponds to the scenario plotted in yellow in Figure 2.8, and has extremely negative cumulative emissions for much of the scenario. Grey shading shows 5%-95% relative likelihood of other parameter combinations, plotted in order of increasing likelihood, showing the uncertainty in peak warming arising from parameter uncertainty in the simple model, as explained in section 2.6.2. The vertical scale bars show the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively. Adapted from Allen *et al.* (2009c)

net cumulative emissions) in the 22$^{\text{nd}}$ Century, meaning that the carbon cycle has not had time to remove as much carbon from the atmosphere at the point at which peak warming occurs, compared with the other pathways, and hence the peak warming is higher for the given cumulative emissions.

Achieving net negative cumulative emissions by 2100 would require dropping emissions to zero instantaneously and then performing the geoengineering technique known as carbon dioxide removal (CDR) to sequester carbon from the atmosphere for the coming century at a faster rate than we have been emitting carbon over the past century, or not dropping emissions to zero but instead removing even more carbon from the atmosphere.

Despite the extreme nature of this sort of emissions pathway, it was still not more than a few tenths of a °C away from the curve showing the relationship between peak warming and cumulative emission. This suggests that pathways with strong negative emissions have the potential to damage the relationship between cumulative emissions and peak warming, but not substantially.

Figure 4.2 shows that the concept of cumulative emissions is robust to even extreme emissions scenarios which oscillate between large negative and large positive emissions each decade. The figure also shows that the relationship between CO_2-induced peak warming and cumulative emissions begins to break down when cumulative emissions go heavily net negative.

In this plot and the next, I have extended the plot range out to include high cumulative totals, and high levels of warming. I have done this not to suggest that this is physical, but simply to illustrate the behaviour of the box-diffusion model in this region.

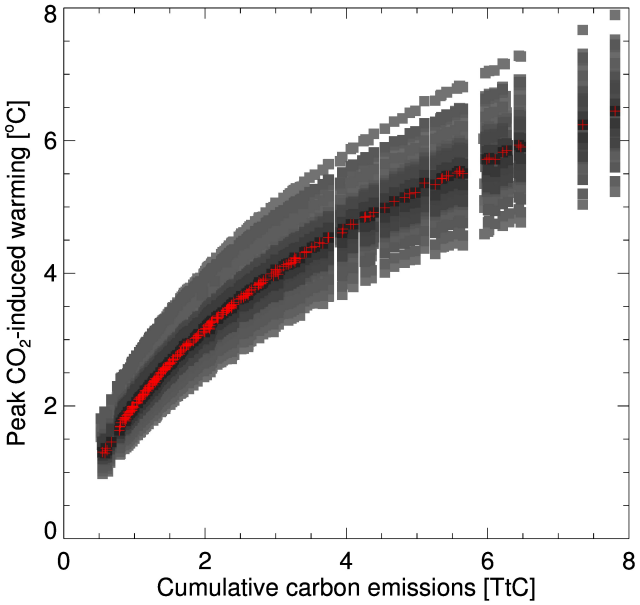
Having outlined the equations used in the model and the emissions time series that are used, and having tested the relationship between cumulative emissions and peak warming under some extreme emissions pathways, I will now explain how we treat uncertainty and determine the parameters to be used in the model.

4.4 The curvature of cumulative emissions vs peak warming

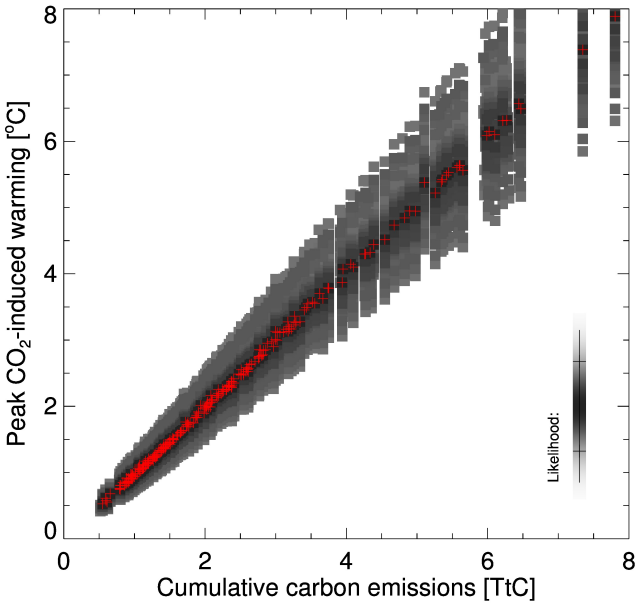
An interesting feature of the tight correlation present in Figures 4.1a and 4.2 is its curvature. Forcing due to CO_2 is proportional to the logarithm of the fractional change in atmospheric CO_2 since the pre-industrial era as discussed in Section 2.4.2 (Ramaswamy *et al.* 2001). If the forcing were linear, the model used in this paper suggests that there would be a more linear relationship between cumulative emissions and peak warming. This is illustrated in Figure 4.3. Panel a reproduces Figure 4.2a but with only the extreme emissions pathways, while panel b is as panel a but with the logarithmic term in Equation 2.17 replaced with a linear term.

Panel b of Figure 4.3 looks linear and indeed shows a proportionality between cumulative emission and peak warming. Panel a, which uses a logarithmic CO_2 forcing like that of the real world, does not show this linearity or proportionality. This result appear to contradict that of Matthews *et al.* (2009), which argues that cumulative emissions and warming are not just linear but proportional, however Matthews *et al.* (2009) only investigate emissions pathways with lower cumulative totals. In the supplementary information of Matthews *et al.* (2009) the authors do suggest that this proportionality would not hold at high cumulative totals. Thus here we provide additional evidence for the comment made in the supplementary information of Matthews *et al.* (2009) that this deviation from proportionality is due to the logarithmic nature of the CO_2 forcing.

There are a number of factors that affect the curvature of this curve. Carbon cycle feedbacks explained in Section 2.3.5, and other non-linear feedbacks such as a non-stationary climate sensitivity make this curve more convex, increasing $\frac{d^2T}{d\sum E^2}$, where $\sum E$ is cumulative emissions. Factors such as the logarithmic nature of CO_2 forcing and the T^4 nature of the Stefan-Boltzmann law of blackbody radiation make the curve more concave, decreasing $\frac{d^2T}{d\sum E^2}$.



(a) Logarithmic CO_2 forcing



(b) Linear CO_2 forcing

Figure 4.3: Peak warming plotted against cumulative emissions to the time of peak warming, using the extreme emissions profiles as in Figure 4.2 but with fewer emissions profiles. The logarithmic term in Equation 2.17 was replaced with a linear term when creating Panel b. The grey-shaded plume shows the 17-83% uncertainty on the red points, as shown in the vertical scale bar in Panel b. The vertical scale bar in Panel b shows the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively. Note that the likelihood profiles in the lower plot have not been retuned to observations, but are simply a representation of the same quantification of uncertainty on the model parameters as in the upper plot, but using a linear forcing term rather than a logarithmic one.

4.5 Emissions targets to constrain the rate of warming

Thus far I have only considered constraints on peak levels of global warming. A key objection to using peak warming targets in isolation is that the feasibility and cost of adapting to future climate change may also depend on the rate of change and not just the magnitude, of global warming. In order to determine which factors constrain the maximum rate of CO_2-induced warming, I use the same model experiments as reported in Figure 4.1. I now plot the peak rate of CO_2-induced warming as a function of the emissions metrics, as illustrated in Figure 4.4.

Figure 4.4a shows only a slight correlation, which occurs because the peak rate of emission and the cumulative emissions are not completely independent. Consider two emissions pathways with different peak rates of emission and the same rate of emissions decline after the peak: the pathway with the higher peak will lead to a higher cumulative total, as shown in Figure 4.4. As I cap the maximum rate of emissions decline to 10% per year, higher peak emissions rates will have a bias towards larger cumulative totals, which explains the correlation I observe in Figure 4.4a.

For a given peak rate of warming, and hence for a given peak emissions rate, pathways with a lower cumulative total or lower emissions in a given year must have a faster rate of decline after the peak. This phenomenon explains why all of the grey, orange and light blue diamonds appear to the left of the black, red and dark blue diamonds in Figure 4.4a-d. The grey, orange and light blue diamonds are less visible in Figure 4.4e and Figure 4.4f because the black, red and light blue diamonds have been plotted over the top of them.

In all of the smooth emissions pathways considered, emissions peaked between 2010 and 2050 by construction, and thus cumulative emissions between 2010 and 2050 are reasonably well correlated to peak emissions rate, particularly when I only consider pathways with rates of emissions decline between 0% and 4%. This is indicated by Figure 4.4b, where the black crosses are particularly well correlated. This would be unlikely to hold for a wider range of emission scenarios though.

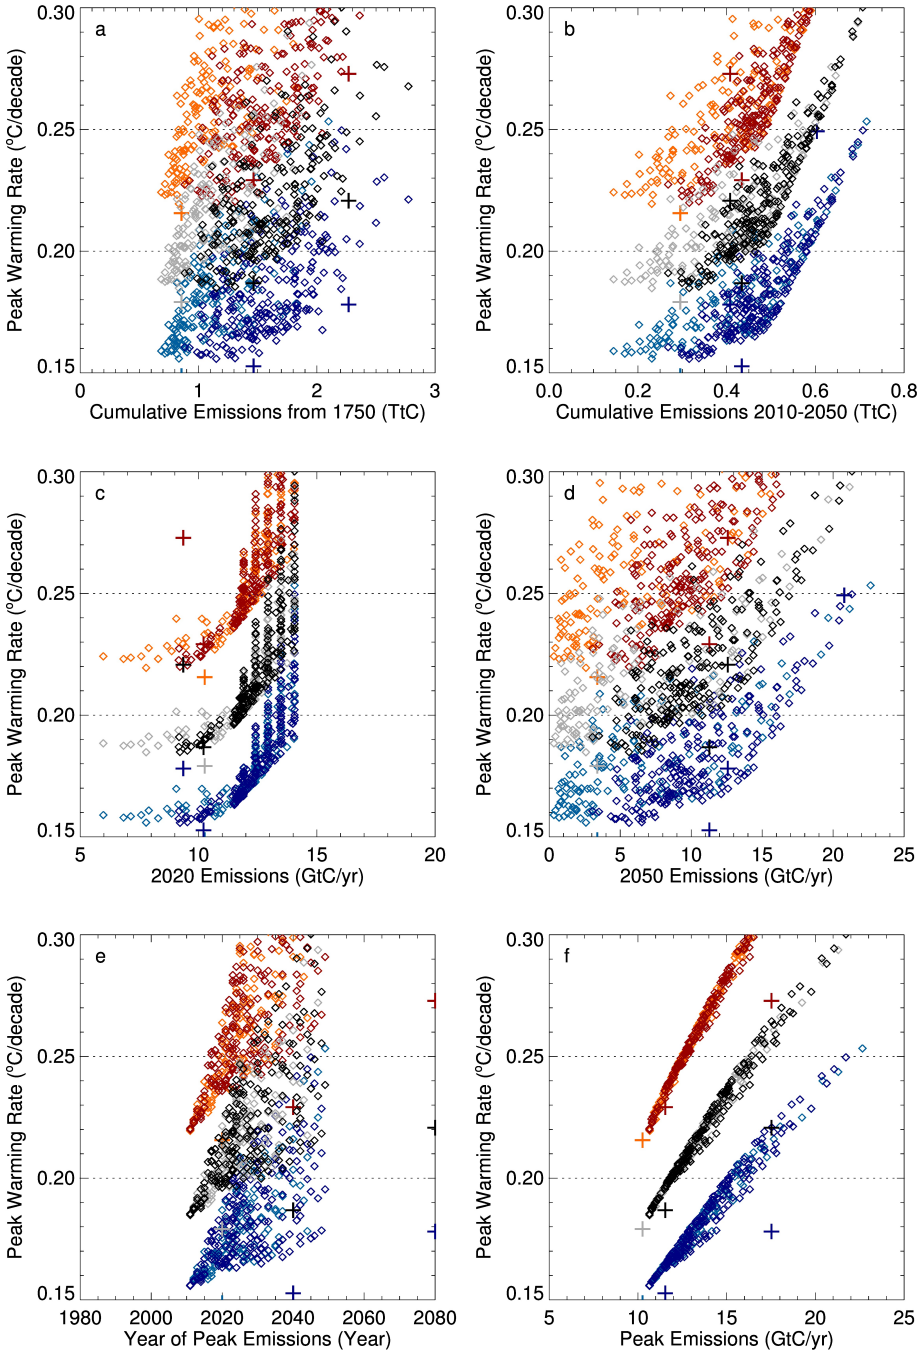


Figure 4.4: Scatter plots showing the relationship between the maximum CO_2-induced rate of warming and various global carbon emission metrics for 395 emission pathways. The x-axis for each panel shows: (a) emissions time-integrated between 1750 and 2500, (b) emissions time-integrated between 2010 and 2050, (c) emissions in the year 2020, (d) emissions in the year 2050, (e) the year in which emissions peak, and (f) the peak or maximum in emissions. Darker diamonds (in black, dark red, and dark blue) indicate emission pathways in which the maximum rate of emissions decline is less than 4%/yr; lighter diamonds (in grey, orange and light blue) indicate the converse. The red and orange symbols are from models at the warmer upper bound of the 17%-83% confidence interval in peak warming. Blue symbols are from models at the cooler lower bound of the 17%-83% confidence interval in peak warming. Grey and black symbols are from models with the best guess peak warming, as discussed in Section 2.6. Crosses represent the RCPs.

Because an emissions peak in the next decade will be heavily constrained by the rate of emission today, Figure 4.4e appears to have some correlation near the present day, which gets worse as I move into the future. The black crosses and the grey diamonds lie in the same region of Figure 4.4e, which suggests that the peak rate of CO_2-induced warming is not heavily affected by the emissions after the emissions peak.

The maximum rate of warming is instead controlled by the peak rate of emission, as indicated in Figure 4.4f. This relationship is not as simple as the relationship between cumulative emissions and peak warming though. Instead this relationship sets an upper limit on the peak warming rate for a given peak emission rate. This is shown by the fact that the upper edge of each of the sets of coloured points in Figure 4.4f is relatively well defined, however the lower edge is more fuzzy and poorly defined, particularly when you include the RCPs.

Any temporary, small, short-term spike in emissions will cause the peak emission rate to increase substantially, but the rate of warming will not increase substantially if the spike was significantly shorter in duration than the climate response timescale and was small enough to not contribute substantially to cumulative emissions (Leith 1975). This is why we can only define an upper limit on the maximum rate of warming for a given maximum rate of emission. This upper limit is determined primarily by the smooth emission profiles, as the smoothest of these will result in the maximum rates of warming possible for given peak emission rates.

The rates of warming for the RCPs are shown in 4.5 in solid line, and the rates of emission are in dotted lines. We see that the relationship between these RCP curves with discontinuous gradients is less strong than Figure 4.4f suggests it is for the smoother curves. Note that we have used only the CO_2 components of the RCPs in this chapter, so the differences between these curves is entirely due to the shape of the CO_2 emissions pathways.

The gradient of the points in Figure 4.4f suggests that for each extra GtC/yr on the peak emissions rate, the best guess maximum rate of warming may increase by up to

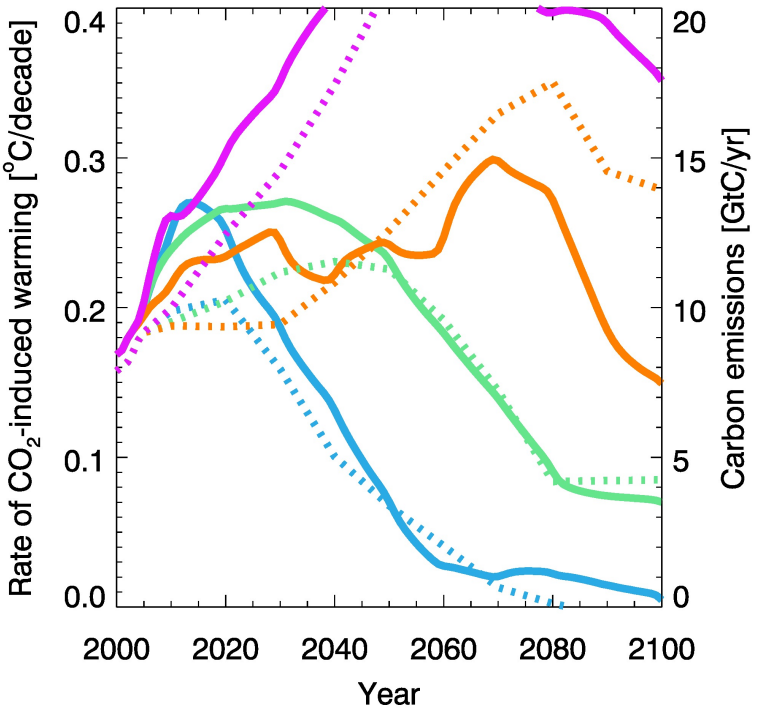


Figure 4.5: Carbon emissions and rate of CO_2-induced warming for the RCPs from the box-diffusion model. Rates of warming are in solid coloured lines (left axis), and carbon emissions are in dotted lines (right axis). The blue curves correspond to RCP-3PD, green to RCP-4.5, orange to RCP-6.0, and pink to RCP-8.5, matching the colours used in Figures 3.4 and 3.3

0.016°C/decade.

Shorter-lived non-CO_2 greenhouse gases, such as methane, which I do not include in this chapter, also influence atmospheric radiative forcing (Forster *et al.* 2007). Although these gases have shorter lifetimes than CO_2 (Forster *et al.* 2007), they still have the potential to influence rates of warming beyond that induced purely by CO_2. I will discuss non-CO_2 GHGs in Chapter 6, but first I will conclude this chapter and discuss the concept of emissions 'floors' in Chapter 5.

4.6 Ongoing and future work

This chapter is about establishing what features of CO_2 emissions controls two indicators of climate change, namely peak warming and rate of warming. There is work to be done in extending this research.

4.6.1 Emissions pathways

In Section 2.5 I introduced the carbon emission pathways used in this chapter. I used a simple algorithm that produces smooth emissions pathways. This algorithm contains a large number of built-in assumptions, such as that once emissions are falling at a certain rate they will always fall at an equal or greater rate of change in the future. Meinshausen *et al.* (2009) uses a more general algorithm for creating emissions scenarios which he introduces in Meinshausen *et al.* (2006), which he calls an *EQW*. I will direct the reader to Meinshausen *et al.* (2006) for a full introduction of the EQW and oversimplify the methodology here as adapting the RCPs or SRES pathways by choosing different fractional rates of change of emissions for different decades while keeping the emissions pathway relatively smooth and plausible. Coding EQWs into the model would allow for the investigation of a far wider range of possible emissions scenarios, and has the potential to make the figures and results in this Chapter 4 more messy and less obvious, but I would not expect it to change the overall conclusions presented here.

4.6.2 Rate of warming and non-CO_2 GHGs

Section 4.5 considers the impact of CO_2 on peak warming, but non-CO_2 GHGs also play a major role in determining the rate of warming. I hypothesise that the rate of warming is controlled by the rate of long-lived climate forcer (LLCF) emissions plus the rate of change with time of short-lived climate forcer (SLCF) emissions. This is essentially the derivative with respect to time of result we reached in Smith *et al.* (2012). This result has to my knowledge never been tested or published and would make an interesting extension of this work.

4.6.3 Constraining other indicators of climate change

This thesis is about avoiding "dangerous" climate change, but peak warming and the rate of warming are just two indicators of climate change. One could extend this work by considering what attributes of emissions control other potential indicators of climate change, from the timing of spring blooms, to crop failure, to species extinction, to just about any other climate impact imaginable. To do this work it would be necessary to couple this simple climate model to impact models. By feeding temperature and CO_2 concentration time series into these impact models and determining impacts we could determine, for example, whether the impact is path dependent, i.e. whether different temperature time series with the same final temperature have the same magnitude of impact in the impact model. We could use this to determine what type of emissions targets we should consider to limit the damage done via this impact. This is similar to the approach taken by Steinacher *et al.* (2013), but would consider more metrics than simply cumulative emissions, CO_2 concentrations and radiative forcing to 2100. Below I give three examples of how this could be done.

4.6.3.1 Ice sheets

Ice sheet melting is thought to be a major contributor to sea level rise (Milne *et al.* 2009). It has been hypothesised that ice melt is dependent on the integral of temperature above

some threshold as a function of time (van den Broeke *et al.* 2010). I hypothesise that this would correspond to the integral of cumulative emissions above some certain threshold with respect to time, but again this has not been tested or published to the best of my knowledge. This simple hypothesis could be tested in the box-diffusion model, though its response to falling temperatures has not been optimised so I would not really trust its results in this regime. It would be more interesting to test this assumption in a model that more closely emulates the C^4MIP ensemble.

This hypothesis could be tested further by coupling a simple carbon cycle-climate model to an ice sheet model. This would allow us to test different assumptions in the ice sheet model and see how they affect emissions targets for avoiding, say, the melting of the greenland ice sheet. However it is unclear whether ice sheet models are currently accurate enough for any conclusions drawn from this exercise not to be dwarfed by model uncertainty (Flato & Marotzke 2013).

4.6.3.2 Extreme weather

We are now able to investigate how climate change is affecting extreme weather using large ensembles of climate models (Pall *et al.* 2010). We could use a platform like climateprediction.net (Stainforth *et al.* 2005) to investigate how future emissions impact extreme weather by running a large scenario and perturbed physics ensemble to establish whether, for example, extreme weather is dependent on emissions path, or purely on the magnitude of total emissions. Climateprediction.net does not have a coupled carbon-cycle, so we would need to generate concentration pathways from a simple carbon cycle like the box-diffusion model, or take a set of concentration pathways, and back out the associated emissions via an inverse calculation.

4.6.3.3 Damages

Ultimately what most people care about is not simply limiting impact on spring bloom times, or ice sheets, or extreme weather, but rather limiting damages, where damages is

defined as some function of all of these things, plus economic damages and damages to human welfare. IAMs have been produced to attempt to estimate damages due to climate change, and we could use one of these to determine which emissions metric (e.g. rate, cumulative total, etc.) is the best indicator of climate damages in different IAMs. As with the other experiments suggested here, this would be very simplistic, but it could help us determine whether it is most import to limit, say the rate of warming or the peak warming. Much work has already been done with IAMs, and the result would probably be a complex interplay of a number of factors, but if a simple result did emerge from a wide array of IAMs it could be useful in informing policy decision-making (Fankhauser *et al.* 2001; Klein & Huq 2007).

4.7 Discussion

Allen *et al.* (2009c), Meinshausen *et al.* (2009), Matthews *et al.* (2009) and Zickfeld *et al.* (2009) showed that cumulative CO_2 emissions are a strong predictor of peak warming. Meinshausen *et al.* (2009) use a metric of cumulative emissions between 2010 and 2050, which I show is a less good indicator of peak warming than cumulative emissions to 2500 which was suggested by Allen *et al.* (2009c). In the next chapter I suggest an even better timeframe to integrate cumulative emissions than either of the suggestions made here.

Huntingford *et al.* (2012), Rogelj *et al.* (2011b) and Rogelj *et al.* (2012b) suggest that warming can be constrained using both 2020 and 2050 emissions. However for a given 2020 emission rate and a given 2050 emission rate, if we also impose a smoothness constraint then we are effectively constraining cumulative emissions.

In Bowerman *et al.* (2011), I consider whether peak emissions provide a constraint on the rate of CO_2-induced warming, and conclude that it does. Here, by considering a wider range of emission pathways, I show that peak CO_2 emissions only provides a constraint on the maximum possible rate of warming. Because short-term spikes in emissions do not contribute substantially to the rate of warming, provided they are much shorter than the

climate response timescale, we can only constrain the maximum possible rate of warming. So we could say, for example, that if we constrain the peak emission rate to 11.5 GtC/yr, then we will have constrained the rate of warming to at most 0.2°C per decade, and possibly less.

The analysis presented here breaks down when we consider some aspects of geoengineering. For example, large negative emissions in the form of CDR between the time of peak warming, and the end of the cumulative emissions integration period breaks the correlation between cumulative emissions and peak CO_2-induced warming.

The findings presented here suggest that we could form climate policy based on limiting cumulative emissions of long-lived GHGs in order to limit peak warming, and then also constraining the peak rate of emission in order to constrain the peak rate of warming. This policy would also need to consider non-CO_2 GHGs. I extended this work to include non-CO_2 GHGs in Smith *et al.* (2012), which I discuss briefly in Chapter 7.

In the next chapter I extend this analysis in a different direction from those suggested above. All of the emissions profiles that I have considered thus far assume that the rate of emission will eventually reach zero. In Chapter 5 I introduce emissions pathways that do not reach zero and consider their implications.

Chapter 5

Emissions floors

5.1 Introducing emissions floors

In this chapter I have developed a new set of emission pathways, extending those described previously. These pathways have emissions floors to represent the emissions that are potentially technologically, economically, or politically unfeasible to mitigate (Anderson 2008). Although these are expressed here in terms of CO_2 alone, they could also be assumed to include the impact of the other principle long-lived anthropogenic RF agent, N_2O (Smith *et al.* 2012).

In this analysis I use two types of emissions floors: a hard floor F_H, and a decaying floor, F_D. These two floors take the forms

$$F_H \geq A \tag{5.1}$$

$$F_D \geq B \exp\left(-\frac{t - t_{2050}}{\tau}\right) \tag{5.2}$$

where A and B are constants with units of gigatonnes of carbon per year (GtC yr^{-1}) as given in the legend in Figure 5.1 and represent the size of the emissions floor in the year 2050 ($t = t_{2050}$), and τ is a time constant set to 200 years. Emissions floors are caps below which emissions are not able to fall, so for all t where $t > t_0$ I take whichever is the larger of $E_a(t)$ and $F(t)$ to be our emissions pathway. I consider five alternative emissions floors, (1) no floor, (2) low hard floor, (3) high hard floor, (4) low decaying floor, (5) high decaying floor, which could apply to each of the possible pathways described in

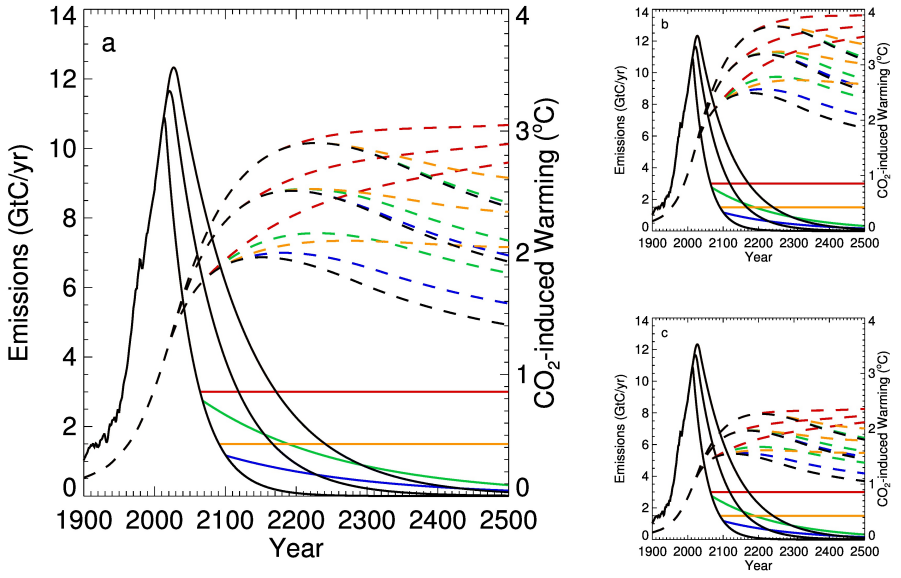


Figure 5.1: Fifteen emission pathways and their resulting temperature trajectories. The emission pathways are in solid lines, and can be read off the left axis, while the temperature trajectories are in dashed lines, and can be read off the right axis. The three parameter combinations used in the box-diffusion model to simulate these responses correspond to the (a) best guess, (b) upper, and (c) lower bound on the 17%-83% confidence interval on expected warming. The panels are identical other than the choice of parameters used in the box-diffusion model. The 15 emission trajectories in each panel are created by combining three possible pathways, shown here in black, with five possible emissions floors, shown here in coloured solid lines, as outlined in Section 5.1. The upper, middle and lower plumes of overlapping coloured dashed temperature trajectories correspond to the black emission profiles that peak the highest, the second highest and the lowest, respectively. The three emission pathways in red solid line with the highest constant emissions floors have red dashed resultant temperature trajectories. The same correspondence applies for the other colours of emissions floors and their resultant temperature trajectories. The upper, middle and lower black curves have cumulative totals to 2500 of 2, 1.5 and 1 TtC, respectively.

Chapter 2. I do not use all of these possible pathways, but rather pick a random subset of a few thousand pathways to investigate with the box-diffusion model outlined in Chapter 2. Fifteen of these pathways are plotted in Figure 5.1, alongside their resulting warming trajectories as simulated by the simple box-diffusion model.

5.2 Emission floors and cumulative emissions

In Figure 5.1 I calculated the warming trajectories not only for emissions pathways with zero-emissions floors, but also for pathways with non-zero floors. I show in Figure 4.1 that cumulative emissions to the time of peak warming are tightly correlated with peak CO_2-induced warming for the case with no emissions floors, and in this chapter I investigate whether emissions floors affect this correlation. Figure 5.2 shows the impact of emissions floors on different cumulative emissions metrics, and each of the panels has the same form as Figure 4.1a. I have plotted most likely peak temperatures as a function of four different cumulative emissions metrics: year 1750 to year 2500 (panel a), year 1750 to the time at which peak warming occurs (panel b), 1750 to year 2100 (panel c), and year 1750 to year 2200 (panel d).

In Figure 5.2a, I can see that pathways with larger emissions floors are shifted to higher cumulative totals. This occurs because the cumulative totals include contributions for portions of the emissions floor that are emitted after the time of peak warming, which can have no effect on peak warming, as illustrated by the green curves in Figure 5.1.

As in Figure 4.1, we see that RCP-3PD, represented by the left-most grey symbols, does not lie on the other curves in 5.2a and d because it has negative emissions between the time of peak temperature and years 2500 and 2200 respectively. In Figure 5.2b, cumulative emissions are integrated to the time of peak warming, and so RCP-3PD does lie on the curve with the other pathways.

I find that in the box-diffusion model if a hard, or non-varying emissions floor becomes too large, then the emissions cannot balance the natural processes that remove carbon from

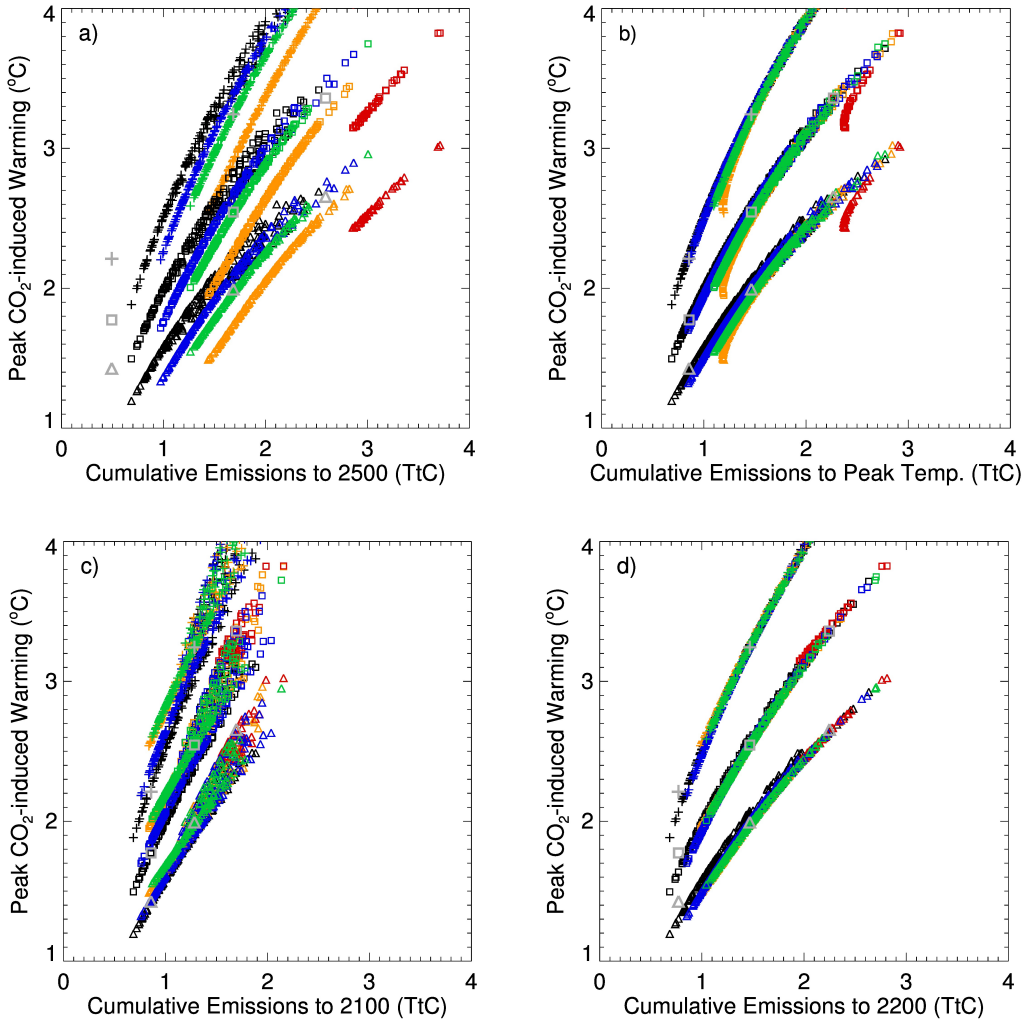


Figure 5.2: Most likely peak warming as a function of cumulative emissions for different emissions floors in the box diffusion model with three different parameter sets. The type of cumulative emission metric varies between the plots: cumulative emissions to (a) 2500, (b) the time of peak warming, (c) 2100, and (d) 2200. The crosses are from the warmer upper bound of the 17%-83% confidence interval on the expected peak CO_2-induced warming, while the squares are from the best guess, and the triangles are from the cooler lower bound. The emissions floors used here are the same as those in Figure 5.1, and use the same colour code. The black symbols represent pathways in which no floor is present, so emissions are allowed to fall to zero. The yellow and red symbols indicate pathways in which a hard floor is set at 1.5 or 3 GtC/yr; in these pathways, emissions are unable to fall below the floor and so remain at these values indefinitely. The blue and green symbols are pathways with an exponentially decreasing emissions floor, which has a decay time of 200 years. The emission pathways represented by the blue symbols pass through 1.5 GtC/yr in the year 2050, while the emissions pathways represented by the green symbols pass through 3 GtC/yr in that year. The grey symbols represent the response to the RCPs.

the atmosphere. At present the precise value of the emissions when the floor becomes too large is uncertain, although I highlight that it may be model specific. This is illustrated by the red curves in Figure 5.1. Consequently, for large hard emissions floors, atmospheric levels of CO_2 continue to rise throughout our 750 year simulation, and are still increasing at the end of the of the experiment, along with associated levels of mean global warming. Extending this analysis to include pathways with cumulative emissions of more than 3GtC, a resultant warming of more than 3-4°C, or cases in which temperatures fail to peak by 2500 would be possible in principle, but would take us outside the range of pathways for which such a simple model is appropriate (see Section 3.7.1). Hence I have not plotted cases where temperatures do not peak by 2500 in Figure 5.2 or Figure 5.3, since I am unable to project when they would peak. All pathways with no floor, and all pathways with a decaying floor, have peaked by 2500, however some pathways with a hard emissions floor have not.

The observational constraints are much more effective in constraining the short and medium-term response of the climate-carbon-cycle system than they are at constraining the multi-century response (Roe & Baker 2007; Allen & Frame 2007). Hence it is inherently hard to determine whether, after a significant injection of CO_2 into the atmosphere, certain emissions floors will cause temperatures to stabilise, decline, or continue rising. Additionally, when considering an emissions floor, it could be argued temperatures will be rising or falling so slowly that their trend could be reversed by actively intervening in the carbon cycle, or simply reducing or increasing emissions by an amount substantially smaller than has already been achieved in reaching that floor. The maximum rate of increase of CO_2 concentrations beyond 2200 associated with the emissions pathways in Figure 5.1 is 28 parts per million by volume per century (ppmv/century). The rate of associated warming shown in Figure 5.1 beyond 2200 is at most 0.27°C/century. In contrast the associated maximum rates in 2100 are a concentration rise of 99 ppmv/century and a warming of 0.88°C/century. It could be argued that a society capable of achieving the kind of rates of emissions reduc-

tion in the year 2100 that are assumed under these pathways would almost certainly be able to convert a static emission floor into a decaying one if it were necessary to do so. Hence it could be argued that the scenario of very rapid reductions followed by a completely stable (hard) floor in CO_2 emissions is unlikely to occur, although the role of N_2O in very long-term emissions is potentially an issue here. Methane could also influence temperatures over the very long term, as could anthropogenic aerosols and other short-lived forcing agents. However, any agent with a lifetime shorter than the response time of the physical climate system must be treated in a very different way to CO_2, since such emissions do not effectively accumulate in the atmosphere over multi-decadal timescales as discussed in Chapter 7.

Figure 5.1 shows that the size and types of emissions floors determine how temperatures will behave after they peak. Several studies have suggested that near-zero emissions are required to stabilise temperatures (Matthews & Caldeira 2008; Plattner *et al.* 2008; Zickfeld *et al.* 2009; Solomon *et al.* 2009). Our simple model's simulations suggest that temperatures will peak then fall slowly under near-zero emissions (Figure 5.1), but this result is acutely sensitive to model structure. At present there appears to be insufficient understanding and suitable observations of the carbon cycle to constrain behaviour during the regime when temperatures decline. In light of this lack of observational constraints, I do not feel confident in relying upon the simple models simulations long after the time at which temperatures peak, as discussed in Section 3.7.1.

Figure 5.2b shows peak warming plotted against cumulative emissions integrated between the year 1750 and the time of peak warming. The correlation in this figure is much better than that in Figure 5.2a. Figure 5.2b shows that a decaying emissions floor does not significantly alter the shape of the relationship between cumulative emissions and peak temperature, as the peak warming is still a function of the cumulative emissions. Emissions floors do, however, affect the lower ends of the curves with low values of cumulative emissions. Consider two emissions pathways, both with a cumulative total of 1TtC, but one

with a decaying emissions floor, and one with no emissions floor: the pathway without an emissions floor will cause a temperature peak earlier than the pathway with the decaying floor, as the emissions floor causes emissions to be emitted over a longer time period. Consequently, in the case with an emissions floor, there will have been more time for carbon to be removed from the atmosphere, presumably resulting in slightly lower atmospheric concentrations at the end of our simulation period than in the no floor case. As forcing is a function of atmospheric concentrations, the case with no emissions floor and higher CO_2 concentrations will result in a higher peak temperature.

I observe this phenomenon in Figure 5.1a by comparing the lowest green and yellow emission pathways and temperature trajectories. The yellow emissions pathway has a higher cumulative total than the green one, when integrated to the time when temperatures peak. Despite this higher cumulative total, the green curve has a higher peak warming than the yellow curve because its emissions are put into the atmosphere over a shorter time period. It is this phenomenon that causes the hard emissions floors to 'peel away' from the soft emissions floors in Figure 5.2b.

In Figure 5.2b, at the upper end of the curve, where cumulative totals are large, the existence of an emissions floor seems to make little difference to the peak temperature. This is because the fraction of the cumulative total that is part of the emissions curve is much larger than the fraction that is in the emissions floor. For the decaying emissions floor in particular, the floor will have decayed to near-zero by the time temperatures peak, as the pathway will reach the floor at a later time than it would have if it had a smaller cumulative total. In general, if the cumulative emissions over the duration of the emissions floor are small compared to the overall emissions, then the floor is not particularly important. If the cumulative emissions over the duration of the floor are a large fraction of the cumulative total, then the level of the floor is a crucial determinant of peak warming. This phenomenon is illustrated in Figure 5.1 by considering the upper yellow and black emissions and temperature curves. The emissions pathway is so large that the yellow emissions floor does

not affect it until 2240, and as a result the yellow and black temperature trajectories are indistinguishable until after temperatures have peaked. This illustrates why emissions floors have less impact on peak warming for pathways with high cumulative totals.

In Figure 5.2c, I see that the correlation between peak warming and cumulative emissions to 2100 is relatively weak. The points furthest to the right of the plot, however, are all black crosses, representing emissions pathways with zero emissions floors. This is because for pathways with zero emissions floors, more of the total cumulative emissions have been emitted by 2100 than for pathways with non-zero emissions floors. In Figure 5.1 all the 15 temperature trajectories are still warming beyond 2100, and all emissions pathways are still emitting beyond 2100. These emissions beyond 2100 are not accounted for in this metric, but will influence peak warming, which accounts for most of the lack of correlation in Figure 5.2c.

The best correlation of all the panels in Figure 5.2 can be observed in Figure 5.2d. This suggests that cumulative emissions, when calculated between 1750 and 2200, are a strong indicator of most likely peak CO_2-induced warming regardless of the type of emissions floor chosen. This is presumably because most of the warming trajectories peak within a few decades of 2200. Those trajectories considered here that do not peak near 2200 have all warmed to within a small fraction of their peak warming by this date, and therefore the emissions emitted in these pathways after 2200 only serve to maintain temperatures, not to induce more warming. This phenomenon is illustrated by the lowest yellow curve in Figure 5.1a, which peaks in 2273, but has warmed to 99% of its peak warming by 2200. This example illustrates how emissions after 2200 have a very small influence on an emission pathways peak temperature, provided the emissions floor is not so high that it prevents temperatures from peaking until beyond 2500. I cannot be confident in the correlation between peak temperatures and cumulative emissions to 2200 for emissions pathways that result in warming of more than 4°C, because peak temperatures will not necessarily occur near the year 2200 in these pathways.

5.3 Likelihoods

In Chapter 4, I confirmed the tight correlation between cumulative emissions and peak CO_2-induced warming, refined in Section 5.2 to consider the effect of non-zero emissions floors. I find that, even with non-zero emissions floors, cumulative emissions, particularly cumulative emissions to the year 2200, correlate well with the resultant peak warming below 4°C.

However, thus far, and in Figures 4.1 and 5.2, our estimates have been of high-, low- and best-guess or most likely response warming, as defined in Section 2.6. In this section I estimate our level of confidence in these results. I re-run our model but with perturbed parameterisations for each ensemble member. For each ensemble member, I can determine a relative likelihood through comparison against our knowledge of the historical record (Pawitan 2001). As explained in Section 2.6.2, models that better reproduced our constraints have higher relative likelihoods.

In Figure 5.3, I do not plot the location of each ensemble member, but instead I plot the outline of the entire ensemble. This allows our likelihood profile to be independent of sampling strategy, provided that I have sufficiently explored parameter space (Pawitan 2001). For each emissions profile within 1% of 1.0, 1.5 or 2.0 TtC cumulative emissions between 1750 and 2200, I calculate a likelihood profile, such that each panel in Figure 5.3 actually contains dozens of likelihood profiles plotted on top of each other.

I note that Figure 5.3c only contains three likelihood profiles, as I consider only three emission pathways with a hard emissions floor and a cumulative total to 2200 of within 1% of 1TtC. Cumulative emissions to 2000 are approximately 0.5 TtC, and a 1.5 GtC/yr emissions floor between 2000 and 2200 has a cumulative total of 0.3 TtC, which leaves only 0.2 TtC remaining if the pathways are to have a cumulative total of 1 TtC. This forces the emissions profile to have a high rate of decline, which could make these profiles socio-economically unfeasible (den Elzen *et al.* 2007).

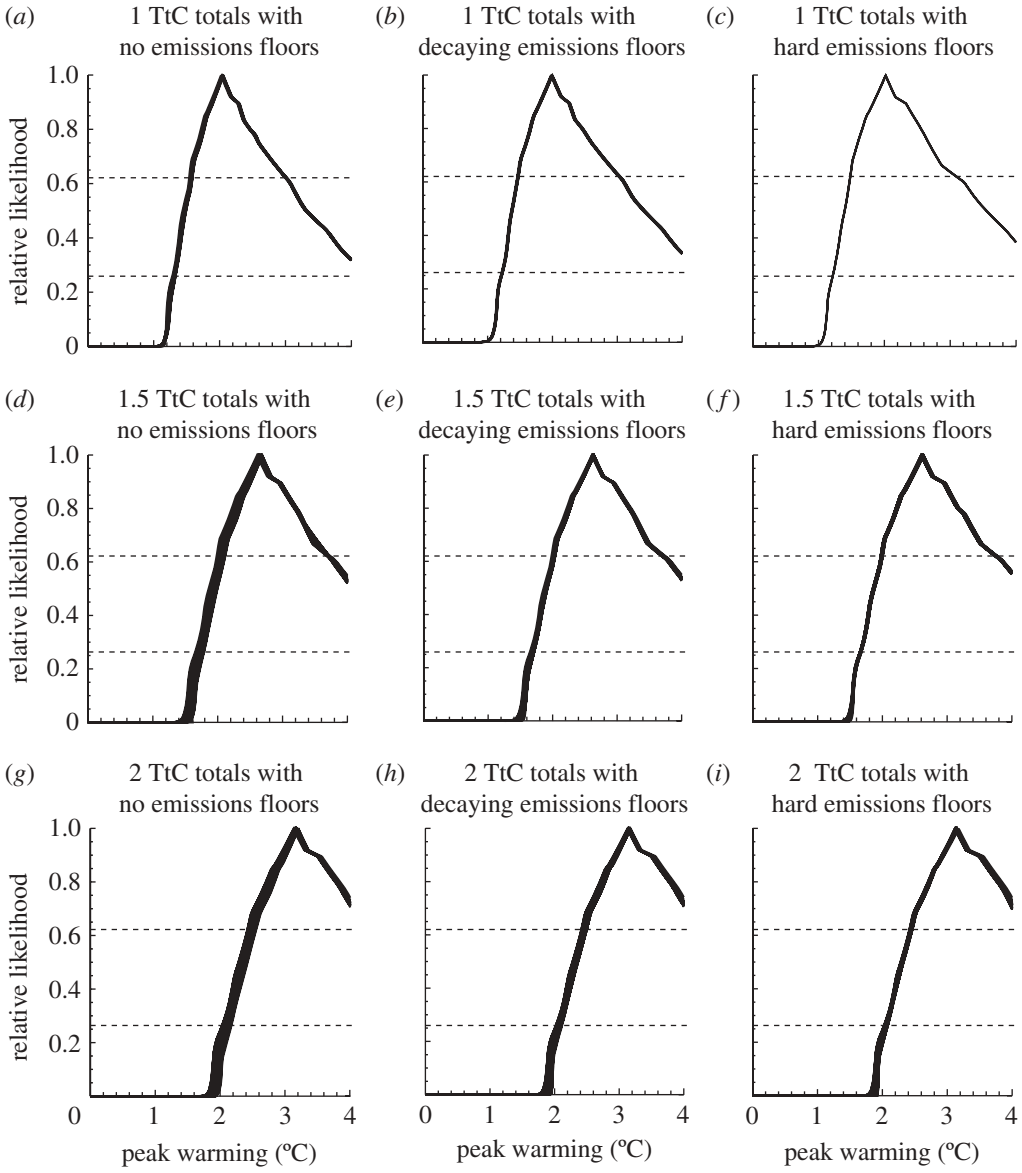


Figure 5.3: Peak warming for different cumulative totals and different emissions floors. These likelihood profiles are produced as outlined in Section 2.6.2. Horizontal dotted lines show thresholds for the 17-83% and 5-95% confidence intervals (Pawitan 2001). Panels (a,d,g) have no emissions floor, so emissions are allowed to fall to zero. Panels (b,e,h) have a decaying (i.e. exponentially decreasing with a 200-year lifetime, passing through 1.5GtC/yr in 2050) emissions floor. Panels (c,f,i) have a 1.5GtC/yr hard emissions floor. In each panel, I plot likelihood profiles over each other for a number of emission pathways described by Equation 2.25 with a cumulative total from 1750 to 2200 within 1% of the stated cumulative total. A sample emission pathway for each of the plots above is given in Figure 5.1, alongside its resultant warming trajectory. The profiles with no emissions floors appear to be drawn thicker only because more emission profiles have been plotted upon one another. I see that the introduction of an emissions floor has little influence on the likelihood profile. Source: Bowerman *et al.* (2011)

All of the likelihood profiles plotted in Figure 5.3 are quite similar, which shows not only that the best-guess peak warming is independent of emission pathway for a given cumulative total, but also that the entire likelihood profile shares this property. I repeat this process for each type of emissions floor so that I can compare likelihood profiles between types. By comparing the likelihood profiles for emission pathways with the same cumulative total but different emissions floors (e.g. the profiles in Figure 5.3d-f), I find that the likelihood profile is unaffected by the type of emissions floor.

5.4 Ongoing and future work

5.4.1 Model dependence

The most striking finding presented in this chapter is how even with emissions floors, the correlation between cumulative emissions and peak warming collapses down to a single line when we integrate to 2200, as illustrated in Figure 5.2. The date of 2200 is probably model dependent, depending on how fast the carbon cycle takes carbon out of the atmosphere at low emission rates. Running this same experiment in other models would be an interesting test of how model dependent this is. This was almost done in a paper that I co-authored, Huntingford *et al.* (2012) using MAGICC. In Figure 3 of Huntingford *et al.* (2012) we plot cumulative emissions to 2500 against the probability of exceeding 2°C for different emissions floors. This shows the same floor dependence as I have in Figure 5.2a, however unfortunately due to time constraints we were unable to reproduce Figure 5.2d to test my hypothesis. Ultimately the integration date at which cumulative emissions correlates well with peak warming is heavily dependent on how much carbon society can emit for the climate-carbon cycle system to be in equilibrium.

5.4.2 Emissions floors in equilibrium

To keep temperatures constant we will have to emit a certain rate of carbon per year, E_{eq}, which is probably close to zero (Friedlingstein *et al.* 2011; Gillett *et al.* 2011; Matthews

& Caldeira 2008; Plattner *et al.* 2008; Weaver *et al.* 2007). In the simple box-diffusion model used here, and in MAGICC, E_{eq} is not constant, is reasonably large, and it is both time-varying and scenario dependent. Several papers using models ranging from EMICs (Matthews & Caldeira 2008; Plattner *et al.* 2008; Solomon *et al.* 2009) to ESMs (Frölicher & Joos 2010; Gillett *et al.* 2011) have suggested that $E_{\mathrm{eq}} \approx 0$ by bringing emissions to zero and showing that temperatures stay constant for several decades, despite gradually falling atmospheric concentrations. The papers explain that this phenomenon is due to the thermal inertia in the oceans. This behaviour is a significant difference between simple models and state-of-the-art models, and suggests that simple models need to be improved to in order to make useful projections of emissions floors given that temperatures in simple models do not stabilise following the cessation of emissions.

5.4.3 The composition of emission floors

Thus far we have only considered emission floors made of CO_2. N_2O and methane (CH_4) are also likely to be a significant components of a potential future emissions floor given their importance in food production. Unless we either dramatically alter or reduce fertiliser use, or see widespread use of genetically modify crops that do not produce much N_2O or CH_4 emissions, or are able to capture or sequester emissions of these gases, food production is likely to make a significant contribution to future emissions. The analysis presented here could be redone to include non-CO_2 GHGs. It would also be interesting to do further analysis into emissions floors, expanding on the work of Anderson & Bows (2011). For example, what types of sectors are likely to contribute to emissions floors and what is currently preventing that sector from reducing its emissions to zero.

5.4.4 Cumulative emissions in a large GCM ensemble

Thus far no-one has used a large scenario or parameter ensemble in a GCM to test the hypotheses outlined in this chapter. This could be done using, for example, climatepre-

diction.net for the climate component (Frame *et al.* 2009), perhaps using a model such as IMOGEN to do the carbon cycle modelling (Huntingford *et al.* 2010).

5.5 Discussion and conclusions

The first use of an emissions floors that I have been able to find in the literature is by Anderson (2008), though the authors do not call this feature an emissions floor until their Anderson & Bows (2011) paper. Since then, the term has been used as a hypothetical minimum level that we are unable to mitigate emissions below (Meinshausen *et al.* 2011c; Huntingford *et al.* 2012; Friedlingstein *et al.* 2011)

In this chapter, and in the Bowerman *et al.* (2011) paper that it is based on, I suggest a new form of emissions floor. Up until now, all of the emissions floors that I have seen in the literature are what I call here a hard or constant floor. I postulate the existence of a future decaying emissions floor. The motivation for this form of floor being that one day society will have mitigated the easier sources of GHG emissions, and will then have to start working on the more challenging sources, which we are only able to mitigate far more slowly.

I use these emissions floors to suggest that year 2200 is a better end date for the integration of cumulative emissions, rather than year 2050 as suggested by Meinshausen *et al.* (2009), year 2100 as suggested by (Anderson 2008) and (Anderson *et al.* 2008) or year 2500 as suggested by Allen *et al.* (2009c). Because of the different impulse response function to CO_2 in the box-diffusion model, as discussed in Section 3.4, the exact date of this integration is likely to be highly model dependent.

I would like to see this result confirmed in models with more complex carbon cycles, as it is likely to be dependent on the long-term response of a model to CO_2 emissions, which varies between models, and which is has the potential to be different in the box-diffusion model compared to EMICs and ESMs as it was not constrained when we tuned the box-diffusion model.

Allen *et al.* (2009c) showed that the entire likelihood profile of the temperature response to the smooth emissions pathways is path independent and depends only on their cumulative emissions in the box-diffusion model. I have extended this work to include emissions floors, and find the same result holds. Note that both of these approaches used smooth emission pathways and the likelihood framework for quantifying uncertainty. In Huntingford *et al.* (2012) Figure 3 we suggest the correlation is not quite as tight as I propose here, but still rather tight, with a span in cumulative probability of 10% or less. In Huntingford *et al.* (2012) we use the MAGICC model, the AVOID emissions pathways which are still smooth, but different from those used here, and use a Bayesian statistical approach.

One way this work could inform current policy targets is for policy makers to view cumulative carbon emissions budgets as spread over, say, four periods: (1) 2010-2020; (2) 2020-2050; (3) 2050-2100 and (4) 2100-2200. Subject to the constraints and caveats outlined above, decision makers have some flexibility in moving emissions from period to period; the important thing for a maximum temperature target is that the overall budget not be exceeded, since this is the primary determinant of peak warming. The inter-period flexibility regarding peak temperature targets ought to be of practical value to policy makers, since it allows them to make informed trade-offs between near-term emissions and emissions in the longer term.

Thus far I have only considered CO_2 emissions and CO_2-induced warming. In the next chapter I will model non-CO_2 GHGs in the simple box-diffusion model and in Chapter 7 I will use them to explore the relationship between SLCFs and LLCFs and their impact on peak warming.

Chapter 6

Modelling non-CO_2 forcing agents

Thus far in this thesis I have only modelled CO_2 and considered its impact on warming. In this chapter I will outline how I can model non-CO_2 forcing agents and then in Chapter 7 explain some of the findings that I have made through considering how CO_2 and non-CO_2 forcing agents interact with global warming.

6.1 Modelling concentrations

As with the rest of the box-diffusion model, I will use the simplest possible treatment of non-CO_2 GHGs. I assume that each gas has a single constant lifetime, following the figures given in the IPCC AR4 (Forster *et al.* 2007). This is a major simplification, however I will show that it produces results that are good enough for the purposes of the work that I am doing.

So we assume

$$\frac{dC_i}{dt} = E_i - \frac{C_i}{\tau_i} \tag{6.1}$$

where C_i is the concentration of GHG i and τ_i is the lifetime of GHG i. E_i is the annual emissions of GHG i converted to have units of concentration change per year where concentration is measured in the same units as C_i as given in Table 6.1. This table shows the gas properties that are needed to model these GHGs using this simple setup, and uses values from Ramaswamy *et al.* (2001) and Forster *et al.* (2007). The gases in Table 6.1 are

the subset of the dozens of potential GHGs that had a RF in 1998, $\Delta F_{1998} \geq 0.01$ Wm^2. The radiative properties of the GHGs will be discussed in Section 6.2.

6.1.1 Comparison with MAGICC

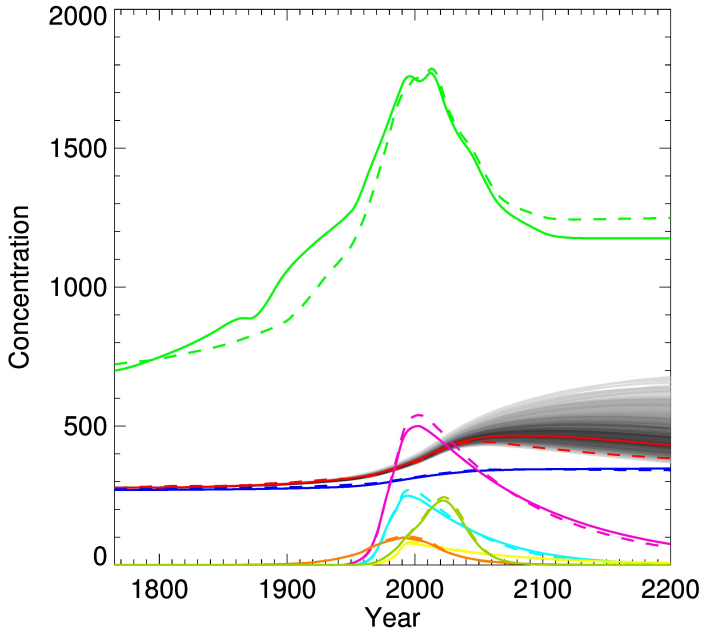
Figure 6.1 shows the concentrations of the eight GHGs as modelled by the box-diffusion model in solid lines. Concentrations are shown for RCP-3PD and RCP-4.5, which are concentration pathways as introduced in Section 2.5. For comparison I also plot the same results as calculated using MAGICC 6 (Meinshausen *et al.* 2011b). MAGICC 6 is classed as another simple climate model, however it is significantly more complex than the box-diffusion model (Meinshausen *et al.* 2011a). The results from MAGICC are qualitatively similar to those from the box-diffusion model, but I will outline the main differences between the two models' results in the sections below. Apart from CO_2, most of the differences between the two models arise because MAGICC has modelled additional complexities. While I could have spent some time in my DPhil adding some of MAGICC's complexities to my model, or converting all my work over to MAGICC, I decided that the simple treatment in the box-diffusion model was good enough for the work I use it for, and opted instead to use the extra time to explore the relationships between SLCFs and LLCFs in Chapter 7. I do use MAGICC to check one of my results in Figure 7.6 which I helped contribute to Smith *et al.* (2012). Next I will outline why the differences in Figure 6.1 arise.

6.1.2 CO_2

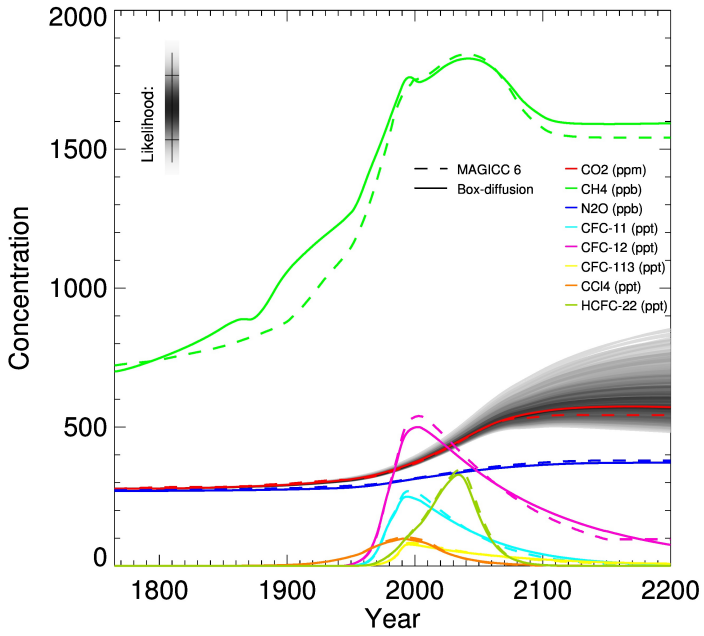
One of the largest discrepancies between the box-diffusion model and MAGICC 6 is in its treatment of CO_2. The MAGICC 6 ocean carbon cycle does not contain explicit diffusion (Meinshausen *et al.* 2011a). This model setup takes carbon out of the atmosphere too quickly when carbon emissions have peaked and fallen to levels significantly lower than they are today, compared with state-of-the-art EMICs and ESMs (Solomon *et al.* 2009; Matthews *et al.* 2009; Matthews & Weaver 2010). As we do not have any observations of carbon emissions peaking and then falling to significantly lower levels, it is difficult to

Trace gas	**Abundance** (Year 1750)	**Abundance** (Year 1998)	**Units of abundance**	**Lifetime** (years)	**RF** (1998) (Wm^2)	**Simplified expression for RF** (Wm^2)
CO_2	278	365	ppmv	N/A	1.46	$\Delta F = 4.841\ln(C/C_0) + 0.0906(\sqrt{C} - \sqrt{C_0})$ $\Delta F = 5.35\ln(C/C_0)$
CH_4	700	1745	parts per billion by volume (ppbv)	12	0.48	$\Delta F = 0.036(\sqrt{M} - \sqrt{M_0}) - (\mathrm{f}(M, N_0) - \mathrm{f}(M_0, N_0))$
N_2O	270	314	ppbv	114	0.15	$\Delta F = 0.12(\sqrt{N} - \sqrt{N_0}) - (\mathrm{f}(M_0, N) - \mathrm{f}(M_0, N_0))$
CFC-11	0	268	parts per trillion by volume (pptv)	45	0.07	$\Delta F = 0.25 \times 10^{-3}(X - X_0)$
CFC-12	0	533	pptv	100	0.17	$\Delta F = 0.32 \times 10^{-3}(X - X_0)$
CFC-113	0	84	pptv	85	0.03	$\Delta F = 0.3 \times 10^{-3}(X - X_0)$
CCl_4	0	102	pptv	26	0.01	$\Delta F = 0.13 \times 10^{-3}(X - X_0)$
HCFC-22	0	132	pptv	12	0.03	$\Delta F = 0.2 \times 10^{-3}(X - X_0)$

Table 6.1: Chemical and radiative properties of the eight greenhouse gases that can be modelled explicitly in the box-diffusion model taken from the IPCC reports. Note that the values used in the box-diffusion model are not all the same as those given here, and are shown in Tables 6.2 and 6.3. C, M, N and X are concentrations of CO_2, CH_4, N_2O and (H)CFCs respectively in the relevant units of abundance. X corresponds to the concentration of (H)CFC in the same row in the table. Subscript 0's represent the preindustrial concentration of a species.
$\mathrm{f}(M, N) = 0.47\ln[1 + 2.01 \times 10^{-5}(MN)^{0.75} + 5.31 \times 10^{-15}M(MN)^{1.52}$.
Source: This table is an amalgamation of tables 6.1 and 6.2 in Ramaswamy *et al.* (2001) and table 2.14 in Forster *et al.* (2007).



(a) Concentrations under RCP-3PD



(b) Concentrations under RCP-4.5

Figure 6.1: Concentrations of our eight modelled GHGs as projected by MAGICC 6 and the box-diffusion model under (a) RCP-3PD and (b) RCP-4.5. Results from MAGICC 6 are plotted in dashed lines, and results from the box-diffusion model are plotted in solid lines. All of the gases in the box-diffusion model are given the simple exponential treatment introduced in Section 6.1 except for CO_2, which is modelled as explained in Section 2.3. The grey shaded plume shows the 5-95% confidence interval on the concentration of CO_2, as explained in Section 2.6. The vertical scale bars show the corresponding likelihood profile for a normally distributed quantity, with black line and tick marks showing the 5-95% and 17-83% confidence intervals respectively.

know how the carbon cycle will behave in this regime. After discussion with a number of experts, in the box-diffusion we chose to include an explicit diffusion term and tune the model to C^{4}MIP (Friedlingstein *et al.* 2006), which produces a model that takes carbon out of the atmosphere more slowly than MAGICC, at least until temperatures peak. This is what accounts for the significant discrepancy between the box-diffusion model and MAGICC, particularly in lower emissions pathways such as RCP-3PD. The form of the carbon cycle model, and its pros and cons, are discussed in more detail in Chapters 2 and 3.

6.1.3 N_2O and the halogenated gases

MAGICC assumes variable lifetimes for each of the other GHGs in Table 6.1, due to a projected increase in the Brewer-Dobson circulation Butchart & Scaife (2001) and due to changes in OH concentration (Meinshausen *et al.* 2011a). However in the exponential treatment used here, I assume constant lifetimes on these gases as given in Table 6.1. These differences account for the small discrepancies between the models' treatments of these gases shown in Figure 6.1.

6.1.3.1 Limits to this approach

Modelling N_2O using a single constant lifetime is a relatively good approximation, as the lifetime of N_2O does not change dramatically. According to Meinshausen *et al.* (2011a) the feedback of N_2O concentration on its own lifetime accounts for a change in lifetime of less than 1% today compared to the pre-industrial lifetime of N_2O. There is approximately a three-year delay between the emission of N_2O at the surface, before it can be transported to the stratosphere where it can decay (Meinshausen *et al.* 2011a). Relative to the 121-year lifetime of N_2O this delay is short, and so we can ignore it for our purposes.

The primary influences on the lifetime of the halogenated gases ((H)CFCs) are changes in OH concentrations which are discussed in Section 6.1.4.1, and the projected changes in the Brewer-Dobson circulation, which affects stratospheric sinks . The Dobson-Brewer circulation is projected to decrease stratospheric lifetimes of the (H)CFCs by 15% per degree

of global mean surface temperature warming (Butchart & Scaife 2001). Given the size of the (H)CFCs forcings, and the size of the discrepancies as shown in Figure 6.1, I am happy to use this simple single-lifetime model of each of the (H)CFCs. I am using this model to compare the impacts of short- and long-lived forcing agents, and the (H)CFCs do not play a large role in this investigation in Chapter 7, and so adding extra complexity to the model of (H)CFCs will not affect the conclusions presented there.

6.1.4 Modelling methane concentrations

The methane results from the two models in Figure 6.1 appear rather different at first sight. The CH_4 model in MAGICC is significantly more complicated than the the simple exponential treatment used in the box-diffusion model. MAGICC incorporates atmospheric chemistry including a temperature dependence of lifetime, the feedback of CH_4 on tropospheric OH and methane's own lifetime, and other effects (Meinshausen *et al.* 2011a). This accounts for much of the discrepancy between the two results.

6.1.4.1 Global methane cycle

Figure 6.2 shows a schematic of the global methane cycle. The first thing to notice in this figure is how large the uncertainty on the numbers is. Compared with the single central estimates given in the schematic of the global carbon cycle in Figure 2.1, the numbers given here do not contain a central estimate but rather a range, with many of the upper bounds several times larger than the lower bounds. As shown in the figure, methane is emitted by a number of diverse and dispersed sources, making it difficult to measure.

Methane is largely removed from the atmosphere via reactions with tropospheric OH, tropospheric chlorine, stratospheric OH, and oxidisation in soils. The first of these reactions removes more methane from the atmosphere than the others combined, and thus it is the one with the shortest lifetime. These lifetimes, τ_i, can be combined in a single lifetime $\bar{\tau}$ using

$$\frac{1}{\bar{\tau}} = \sum_i \left(\frac{1}{\tau_i}\right). \tag{6.2}$$

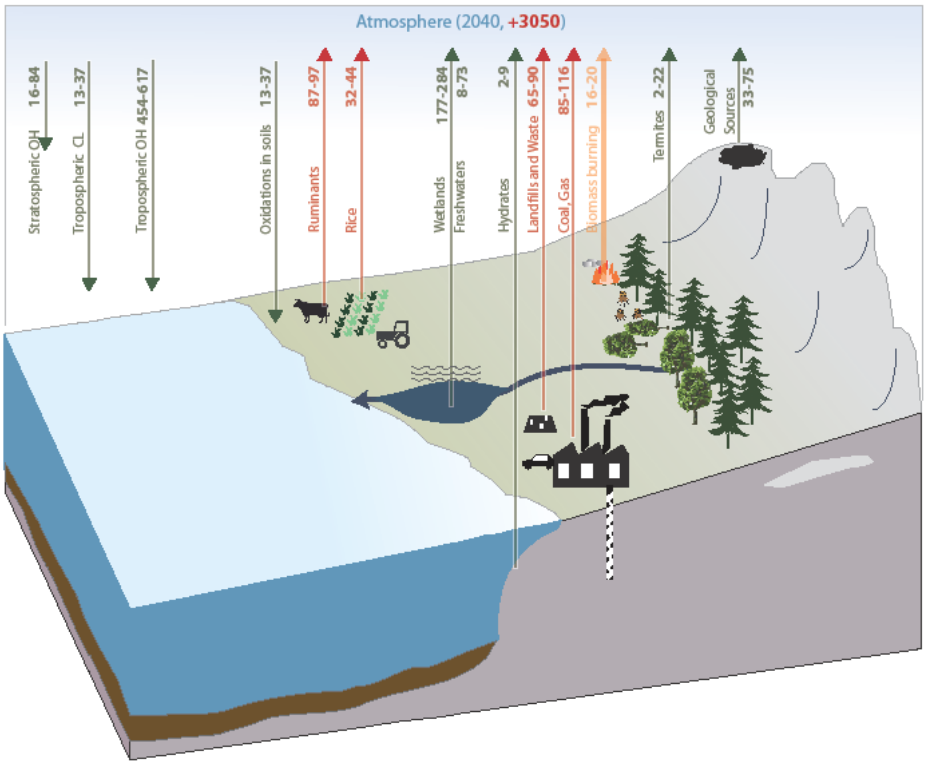


Figure 6.2: A schematic of the global methane cycle. Red lines are anthropogenic fluxes, green arrows represent fluxes due to natural process, and orange lines are combined natural and anthropogenic fluxes (because we are unsure to what extent the emission is anthropogenic rather than natural). Numbers represent methane fluxes in units of $\mathrm{Mt}(CH_4)\mathrm{yr}^{-1}$. Source: Ciais & Sabine (2013)

In the box-diffusion model I assume that $\bar{\tau}$ is a constant, however in reality each τ_i depends on a number of factors, but in particular on the atmospheric concentrations of OH. The concentration of OH depends in turn on the concentrations of NO_x, CO, volatile organic compounds (VOCs) and methane, meaning that in practise the evolution of the atmospheric concentrations of each of these gases depends on each other.

6.1.4.2 Inverting methane emissions

One way of analysing the discrepancy between the box-diffusion model's methane concentrations and MAGICC's methane concentrations is to compare the emissions that would be required in each model to reproduce observed historic concentrations. The change in the concentration of methane over time is determined by Equation 6.1. We can solve for E, and hence determine the emissions required in any given year to produce the change in concentrations that occurred. This process is known as 'inverting' emissions from concentrations, and is what has been done to produce Figure 6.3.

The dark blue and purple lines in Figure 6.3 have been inverted from estimates historical concentrations using Equation 6.1. The red line in the figure has been inverted from concentrations using the MAGICC model (Meinshausen *et al.* 2011c). The other lines are all historical estimates of emissions. The light blue line is a top-down estimate of methane emissions over the past three decades, derived from inversions from satellite observations (Ciais & Sabine 2013). The orange line is a bottom-up estimate of methane emissions taken from a combination of field studies and proxy data such as the area under rice paddy cultivation or estimates of cattle head counts (Ciais & Sabine 2013). The green line is an older bottom-up estimate of methane emissions from 1860 onwards from Stern & Kaufmann (1996).

The error on the light blue and orange lines is taken from the IPCC AR5 (Ciais & Sabine 2013). There is no error provided by Stern & Kaufmann (1996), and so I have chosen to apply a fixed error on their result. I chose 72MtC/yr as it is the smallest error that allows the central estimates of the IPCC bottom-up estimates to fall within the error on the Stern &

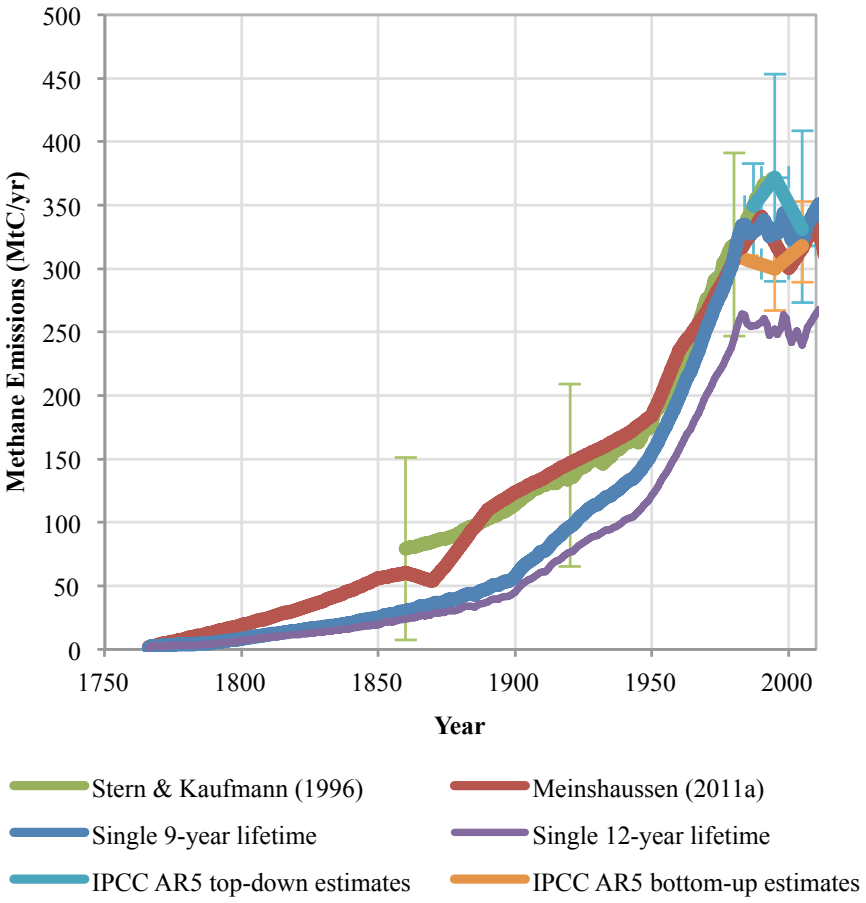


Figure 6.3: Methane emissions from anthropogenic sources. The green line is a bottom up estimate from Stern & Kaufmann (1996), based on proxies including estimates of global cattle numbers and the area of rice paddies. No error on this estimate is given by Stern & Kaufmann (1996), but I have assumed an error of 72MtC/yr, which is the minimum error that allows the central estimate from the bottom-up approach in IPCC AR5 to fall within the error bars of the Stern & Kaufmann (1996) estimate. The red line is from Meinshausen *et al.* (2011c), which has been inverted from concentrations using the simple model MAGICC. The light blue and orange lines are the top-down and bottom-up estimates from Ciais & Sabine (2013), and show their respective uncertainties. The dark blue and thin purple lines are emissions derived from concentrations assuming a single lifetime of 9 and 12 years respectively.

Kaufmann (1996) estimate. This error is similar in magnitude to the error on the top-down global estimates of methane emissions from Ciais & Sabine (2013).

Much of this error derives from the fact that we do not know how much of various emissions are natural rather than anthropogenic, as we are unsure what the natural baseline emission was (Bousquet *et al.* 2006; Etheridge *et al.* 1998b). Additionally, we do not know at which point human influence on the methane cycle became non-negligible, again because we are not sure what fraction of current or historic emissions are natural vs anthropogenic (Ruddiman 2003; Claussen *et al.* 2005; Sapart *et al.* 2013).

Another source of error derives from the fact that we are uncertain about historical OH concentrations, despite the fact that we can measure current OH concentrations (Stern & Kaufmann 1996). Some of these sources of error may result in more uncertainty in absolute terms when emissions were smaller, such as the error on fossil fuel emissions (though even this could plausibly increase historically due to less knowledge about how much fossil fuel extraction was occurring) (Boden *et al.* 2013; Höhne *et al.* 2010). Others sources of error may have been more uncertain in the past in absolute terms than they are today, such as the change in OH concentrations. Consequently I do not think that a fractional error on the anthropogenic emissions is an appropriate estimation of this error, and so I have used an absolute error on the Stern & Kaufmann (1996) estimate.

6.1.4.3 Methane lifetime

As shown in Figure 6.3, a fixed methane lifetime of 12 years does not appear to fit the data historically, but a fixed methane lifetime of 9 years does fit the data within the error bars discussed above. The emissions inverted by the MAGICC model are significantly higher before the 1950s than those inverted using a single fixed lifetime. Indeed, an analysis of the MAGICC inversion using Equation 6.1 suggests that the MAGICC model is using a methane lifetime of only 4 years during the 1800s in order to match the high emissions that have been estimated by Stern & Kaufmann (1996). I have no evidence for the methane lifetime falling this low.

The IPCC central estimate for the lifetime of methane is 12 years, however this number is uncertain. Holmes *et al.* (2013) used three different chemical transport models to estimate the current lifetime of methane, and suggest it is 9 ± 1 years. A 9-year lifetime is a much better fit to the data in this single-lifetime model than a 12-year lifetime, and so I will use a 9-year lifetime here.

I intend to use this single constant lifetime model of methane to compare the impact of short-and long-lived climate forcers. As a result the model does not need to be an accurate representation of methane, it only needs to be good enough that a more accurate representation would not affect the conclusions that I reach in Chapter 7. I have recreated the plots in Chapter 7 with a 12-year lifetime rather than a 9-year lifetime, and there is no qualitative difference in the results presented. I have not included these plots due to the page limit restriction, and because they look very similar to the plots already presented. This finding suggests that I am justified in using a simple representation of methane here.

Having outlined the exponential treatment that we use to model the lifetimes of non-CO_2 GHGs in the box-diffusion model, I will now explain how we model the RF due to non-CO_2 GHGs.

6.2 Modelling radiative forcing

In the atmosphere, RF is determined in part by how much outgoing long-wave radiation escapes to space without being absorbed by the atmosphere. This is determined in part by the radiative properties of the atmosphere, and the spectral properties of the GHGs. Figure 6.4 shows the infrared (IR) spectra of five major GHGs. Note how water vapour (H_2O) is absorbing across much of the IR spectrum, but CO_2, CH_4 and N_2O are absorbing only in certain spectral bands.

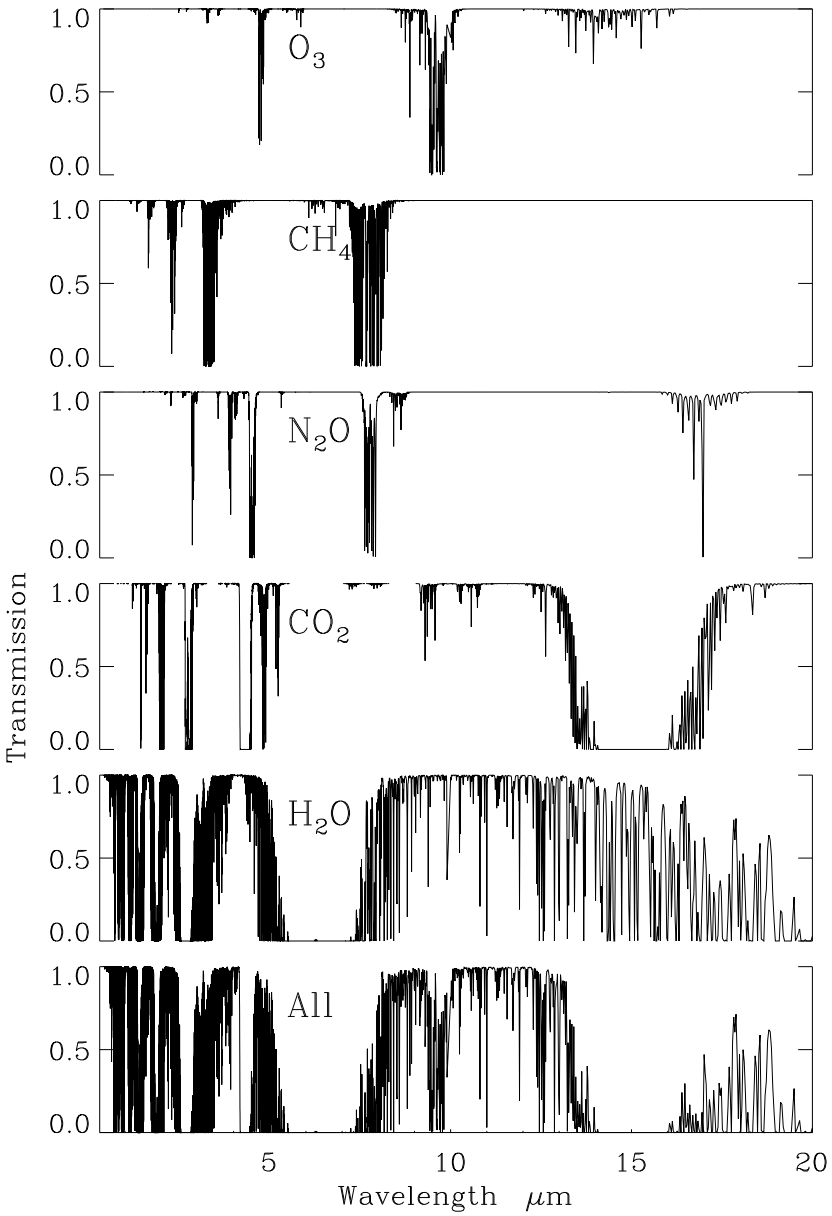


Figure 6.4: The spectra of five GHGs in the infrared region. Transmittance is measured from 0 (no transmittance) to 1 (full transmittance). These spectra are taken from a radiation model which simulates the propagation of radiation through the Earth's atmosphere at different wavelengths. Source: Don Grainger

6.2.1 Methane

Modelling the RF due to adding extra CH_4 to the atmosphere is complicated by the fact that the absorption bands of CH_4 and N_2O overlap near $\lambda = 7.5\mu$m, as illustrated in Figure 6.4. This means that the RF due to CH_4 is dependent on the abundance of N_2O in the atmosphere, and vice versa. In the IPCC Third Assessment Report (TAR), Ramaswamy *et al.* (2001) simplify this relationship to Equation 6.3, as the dependence of the total RF on realistic future changes in N_2O is small compared with the other terms, to

$$\Delta F = 0.036(\sqrt{M} - \sqrt{M_0}) - (\mathrm{f}(M, N_0) - \mathrm{f}(M_0, N_0)) \tag{6.3}$$

where $\mathrm{f}(M, N) = 0.47 \ln [1 + 2.01 \times 10^{-5} (MN)^{0.75} + 5.31 \times 10^{-15} M (MN)^{1.52}]$. M is the atmospheric concentration of CH_4, N is the atmospheric concentration of N_2O, and a subscript 0 indicates a pre-industrial concentration. Note that Ramaswamy *et al.* (2001) linearises F around N_0 in Equation 6.3, such that $F = f(M, M_0, N_0)$ and not a function of N, the current concentrations of N_2O.

6.2.2 A generic radiative forcing equation

Equation 6.3 is a rather complicated equation with a large number of parameters. In keeping with the simplicity of the box-diffusion model, I attempted to simplify it down to a smaller equation with only four free parameters. I wanted this formulation to be generic enough to be applicable to the forcing due to any GHG, but I wanted it to be physically based so that it wouldn't break down as quickly as a purely statistical model if extrapolated beyond its calibration regime. Based on Ramaswamy *et al.* (2001), I chose the form

$$\Delta F = a \ln(X/X_0) + bX + c\sqrt{X} + d. \tag{6.4}$$

where X is the concentration of an arbitrary GHG and a, b, c and d are constants to be found for each GHG. In the following paragraphs I will outline why I chose the four terms in this equation.

Logarithmic term I chose to include a logarithmic term in Equation 6.4 following the logarithmic dependence of CO_2-induced RF introduced and discussed in Section 2.4.2. This will allow me to model GHGs that have logarithmic dependence on concentration.

Linear term In the case of an isolated spectral line we can consider the *weak limit*. In this case we can show via a Taylor expansion of the relevant exponential term that gases which are weakly absorbing, such as the (H)CFCs, have an absorption that is proportional to the mass of the absorber, which is proportional to the concentration of the weakly-absorbing GHG (Wendisch & Yang 2012). This is why Ramaswamy *et al.* (2001) can model the (H)CFCs as having $\Delta F = \alpha(X - X_0)$, where α is some constant. Thus, I include a linear term so that I can model GHGs that are close to the weak limit.

Square root term For a pressure-broadened spectral line we can also consider the *strong limit* in which a line centre is completely absorbing and any increases in mass loading or concentration only affect the wings of the spectral line. In this case we can show via another Taylor expansion of an exponential that in the strong limit the absorption of a gas will be proportional to the square root of the concentration (Wendisch & Yang 2012) so that $\Delta F \simeq \alpha(\sqrt{X} - \sqrt{X_0})$. This is why the first term in Equation 6.3 takes this form. Thus I have included a square root term in Equation 6.4 to account for GHGs that have components of their spectrum that behave close to the strong limit.

Constant term Radiative forcing is defined relative to the pre-industrial era, and so any changes in the concentration of GHGs must be measured relative to pre-industrial. The constant term, d, in Equation 6.4 incorporates terms such as $-\sqrt{X_0}$ required to set the pre-industrial forcing of the GHGs to be zero.

6.2.3 Curve fitting

Having introduced Equation 6.4 and explained the rational behind it, I can now attempt to find the constants a, b, c and d for methane by fitting Equation 6.4 to Equation 6.3 for the range of possible future CH_4 concentrations. This is shown in Figure 6.5a.

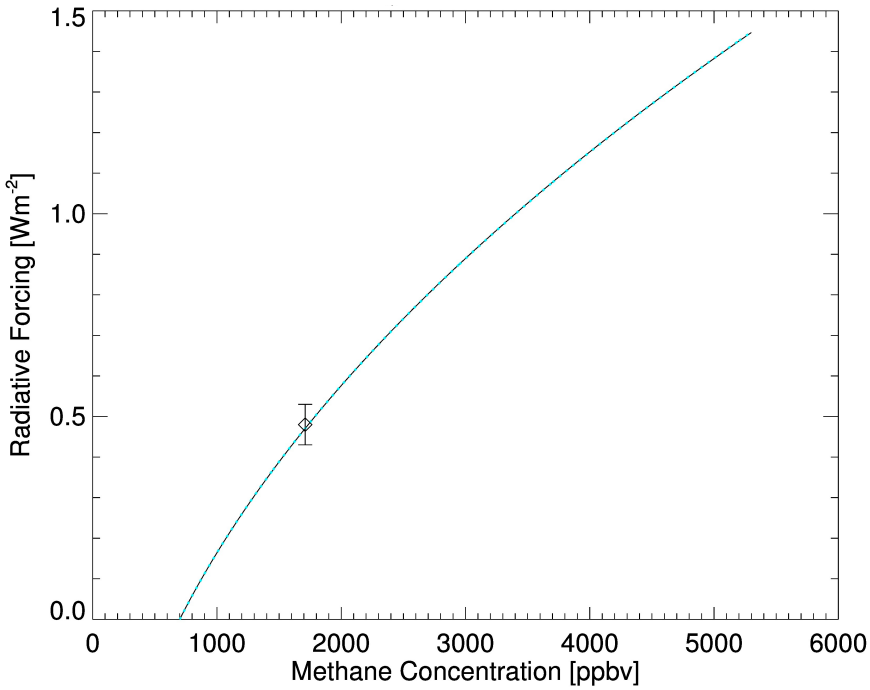
We can repeat this process for N_2O, which has a RF equation similar to that of CH_4 in Equation 6.3,

$$\Delta F = 0.12(\sqrt{N} - \sqrt{N_0}) - (\mathrm{f}(M_0, N) - \mathrm{f}(M_0, N_0)) \tag{6.5}$$

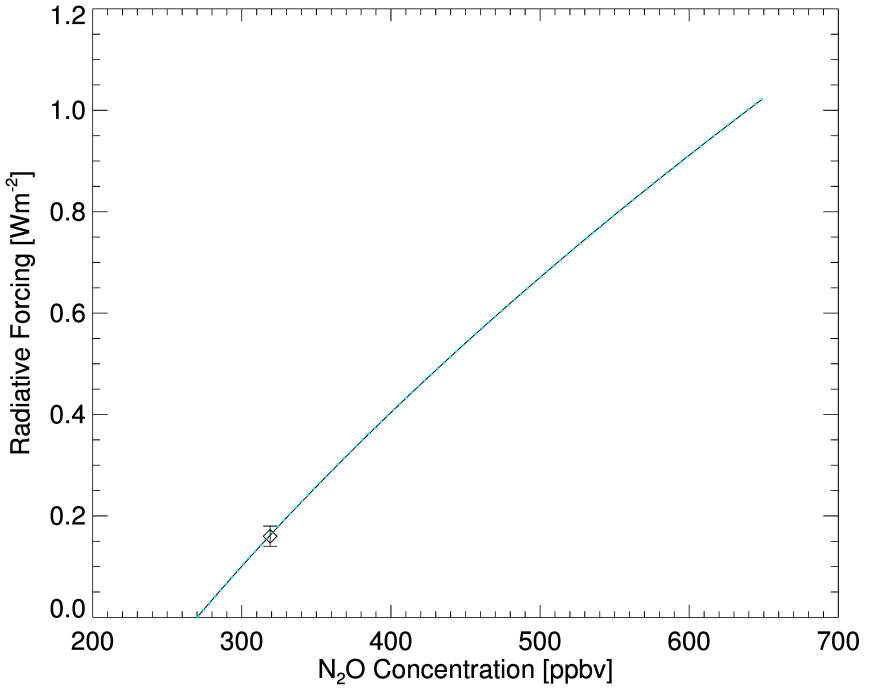
where again $\mathrm{f}(M, N) = 0.47 \ln [1 + 2.01 \times 10^{-5} (MN)^{0.75} + 5.31 \times 10^{-15} M (MN)^{1.52}]$ as in Equation 6.3 (Ramaswamy *et al.* 2001). The results of this curve-fitting exercise are shown in Figure 6.5, which has a dashed curve plotted over a solid curve which is so similar that the two curves are practically indistinguishable. This shows that the curve fitting has been successful in reproducing the IPCC RFs of Equations 6.3 and 6.5, as well as in reproducing present day values of radiative forcing.

Figure 6.5 shows that the forcing due to CH_4 and N_2O can be approximated to a high degree of accuracy by Equation 6.4. This equation contains far fewer parameters than Equations 6.3 and 6.5. Though Equations 6.3 and 6.5 are more physically-based than Equation 6.4, this suggests that, if required, simple climate models could use a far simpler functional form than the equations given by the IPCC to describe the forcing due to CH_4 and N_2O. As a different example, the final term of the expression for $\mathrm{f}(M, N)$ in equation 6.3 contributes less than 1% of the total forcing, and could be ignored and the equation could be retuned to provide another simplified equation for Equations 6.3 and 6.5.

In some ways it is not surprising that I can fit a curve to an equation with four free parameters. As Enrico Fermi attributed John von Neumann as having said, "with four parameters I can fit an elephant and with five I can make him wiggle his trunk" (Dyson 2004).



(a) RF due to CH_4



(b) RF due to N_2O

Figure 6.5: RF due to different atmospheric concentrations of (a) CH_4 and (b) N_2O for two different functional forms of the RF equation. The IPCC functional form in (a) Equation 6.3 (Ramaswamy *et al.* 2001), (b) Equation 6.5 is shown in black, and the fitted function form in Equation 6.4 is shown in light blue dots. In each panel the two curves that have been plotted on top of each other and are the same to within 1% accuracy so they look like they are a single curve whereas actually they are two curves plotted on top of each other. The diamonds and error bars are the Forster *et al.* (2007) estimates.

6.2.4 All greenhouse gases

Now that we have found a generic RF Equation 6.4 that fits CH_4 and N_2O, we can use it to represent all of the other RF equations given in Table 6.1. Thus we can model the RF due to all of the GHGs in Table 6.1 using the generic RF equation I introduced in Equation 6.4. Table 6.2 gives the constants a, b, c and d in Equation 6.4 for each of the GHGs in Table 6.1.

	a	b	c	d
CO_2	4.84	0	0.0906	-1.51
CH_4	0.03	4×10^{-6}	0.0294	-0.78
N_2O	0.42	9×10^{-4}	0.0324	-0.79
CFC-11	0	0.25	0	0
CFC-12	0	0.32	0	0
CFC-113	0	0.3	0	0
CCl_4	0	0.13	0	0
HCFC-22	0	0.2	0	0

Table 6.2: Values for the constants a, b, c and d in Equation 6.4 for eight GHGs. These constants can be used to construct a RF equation for each of the eight GHGs.

Having introduced simple models of the concentrations and RFs of each of the eight GHGs that contribute most to the total RF, I will now explain how I model black carbon.

6.3 Modelling black carbon

In Chapter 7 will be modelling the impact of black carbon (BC) on warming. Black carbon influences RF in two ways: directly via atmospheric absorption, and indirectly via albedo by making snow and ice darker (Boucher & Randall 2013). The direct effect has a lifetime of much less than one year, but the indirect effect has a globally averaged lifetime of slightly longer than one year because it takes time for black carbon on snow to be covered up by fresh snow. The box-diffusion model has a time step of one year so we are unable to model process that occur on timescales shorter than one year. Because these effects have a lifetime

approximately equal to or much shorter than one year, we can approximate

$$\Delta F_{\mathrm{BC}} \approx \alpha C_{\mathrm{BC}} \approx \beta E_{\mathrm{BC}}, \tag{6.6}$$

where ΔF_{BC} is the global-mean, annual-mean RF due to BC including both direct and indirect effects, C_{BC} is the global-mean, annual-mean concentration of BC, E_{BC} is the annual global emission of BC, and α and β are constants to be found.

The unofficial RCPs provide E_{BC} and ΔF_{BC} as simulated using MAGICC and so we can solve for β (Meinshausen *et al.* 2011c). E_{BC} and ΔF_{BC} are plotted in dark blue and red in Figure 6.6. The curve in this Figure has two parts, a historical section in dark blue, which contains more variation, and a predicted future section in red which is made up of a straight line which would intercept the origin if continued. Because we are primarily interested in future changes in BC, we will use the same assumption as was made by Meinshausen *et al.* (2011a) and assume $\Delta F_{\mathrm{BC}} \propto E_{\mathrm{BC}}$ in our treatment of BC. Thus in the box-diffusion model we assume $\Delta F_{\mathrm{BC}} = \beta E_{\mathrm{BC}}$ where ΔF_{BC} has units of Wm^{-2}, β has units of W yrs Mt^{-1} m^{-2}, and E_{BC} has units of Mt yr^{-1}, where Mt means megatonnes.

The unofficial RCPs use present day BC emissions of 8Mt/yr, and a present forcing of 0.3 Wm^{-2}. A recent paper by Bond *et al.* (2013) provides the most comprehensive estimate of BC emissions and forcing to date using primarily top-down estimates, and suggests that most models use BC forcings that are approximately three times to small. Bond *et al.* (2013) suggest that present day emissions are in fact between 8 and 17 Mt/yr, and that radiative forcing is between approximately 0.1 and 2 Wm^{-2}. In Bond *et al.*'s analysis, they uses a central estimate of 14 Mt/yr and 1.1 Wm $^{-2}$, plotted in green in Figure 6.6. In light of this increased estimate, in the box diffusion model I use the unofficial RCP emissions, scaled by a factor of 1.5, as shown in light blue in Figure 6.6 (Meinshausen *et al.* 2011c). I also use β, the conversion factor between BC emissions and forcing from Meinshausen *et al.* (2011c), but scaled up by a factor of two as shown in light blue and purple in Figure 6.6, giving an overall increase in the forcing due to BC of a factor of three, as recommended by Bond *et al.* (2013). Figure 6.6 shows that scaling up the RCPs' BC

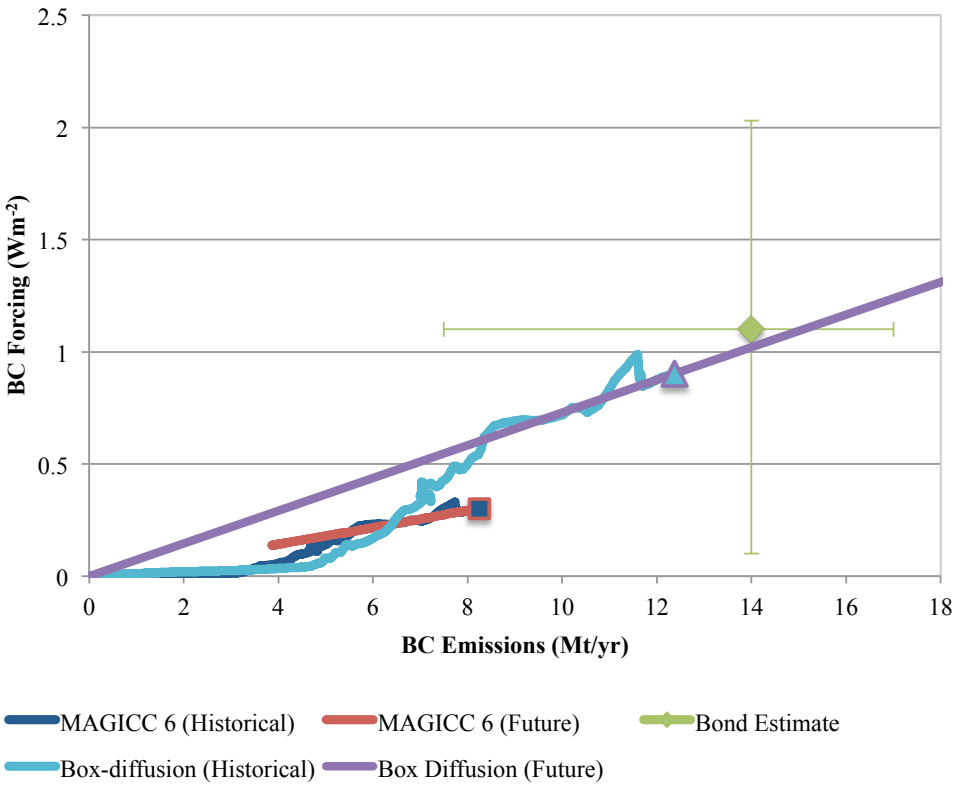


Figure 6.6: The relationship between black carbon emissions and black carbon forcing. The historical relationship from Meinshausen *et al.* (2011c) is in dark blue, and the future relationship used in MAGICC 6 is in red. The present day value used by Meinshausen *et al.* (2011c) is the red and dark blue box. The present day estimate from Bond *et al.* (2013) is in green with associated error bars as discussed in the text. The light blue line obtained by scaling up the Meinshausen *et al.* (2011c) emissions up by a factor of 1.5, and the forcing up by a factor of two, in keeping with the recommendations of Bond *et al.* (2013). This is the relationship between black carbon emissions and forcing used in the box-diffusion model, and the purple and light blue triangle is the present day value used. The purple line is the future relationship between these quantities used in the box-diffusion model.

emissions and forcing in this manner results in a similar result to that suggested by Bond *et al.* (2013).

Tropospheric ozone is modelled in the box-diffusion model in the same way as BC is here, namely by assuming that annually-averaged radiative forcing is directly proportional to annually-averaged emissions. Thus to model the impact of reducing emissions of tropospheric ozone precursors by $X\%$, I simply reduce the radiative forcing due to tropospheric ozone by $X\%$.

6.3.1 Limits to this approach

Despite the fact that we model black carbon as a well-mixed aerosol, its emissions, concentrations, radiative forcing and climate impacts are highly regional (Myhre & Shindell 2013; Zarzycki & Bond 2010). Black carbon emissions, concentration and forcing are all strongest in South-East Asia where biomass burning is greatest, with a large component also coming from Central Africa for similar reasons (Shindell *et al.* 2012).

The forcing due to black carbon is also highly dependent on the albedo of the surface underneath it. High albedo surfaces such as the cryosphere, desert, or clouds produce the largest BC forcing (Bond *et al.* 2013). As much as 50% of BC RF comes from BC above clouds (Zarzycki & Bond 2010) A large component also comes from the Himalayas, with their close proximity to the East Asian sources of emission (Shindell *et al.* 2012).

Because of its regional forcing, BC may also have a regional impact on climate change, particularly via regional phenomenon such as the East Asian monsoon. Black carbon is projected to increase atmospheric forcing (defined as the top of atmosphere (TOA) forcing minus the surface forcing) particularly in parts of Asia and Africa, which can strongly influence regional precipitation patterns (Shindell *et al.* 2012).

The regional impact of BC forcing varies considerably between models, suggesting that the impact of BC is relatively uncertain (Shindell *et al.* 2012). Nonetheless Shindell *et al.* (2012) report the NASA GISS-E2 model shows that reductions in future BC forcing result in a reduction in the shift of the East Asian monsoon relative to scenarios in which BC

emissions reductions are not carried out. They suggest that BC emissions reductions could reduce drought risk in the Sahel and Southern Europe, while reversing the shifting patterns of the South East Asian monsoon. I would be reluctant to read too much in to the result from one model, and instead suggest this is an area deserving of more study, particularly using higher resolution regional models (Flato & Marotzke 2013).

6.4 Uncertainty

In this thesis, I have not attempted to quantify the uncertainty on the box-diffusion model's treatment of non-CO_2 GHGs. This is largely because the lifetime and forcing uncertainties involved are far smaller than the uncertainty in the temperature response. For example the one-σ error on the current lifetime of methane is approximately 10% (Holmes *et al.* 2013), the error on the current forcing due to atmospheric methane concentrations is approximately 10% (Myhre & Shindell 2013), while the upper bound on the one-σ error on the equilibrium warming due to a fixed radiative forcing is 300% larger than the lower bound (Collins *et al.* 2013).

Quantifying the uncertainty in the non-CO_2 GHGs in the likelihood-based framework presented in Section 2.6.2, as has been done in the remainder of this thesis, would be challenging as it would require more observable constraints than we have easily available at the moment.

Each non-CO_2 GHG would require one constraint on its lifetime, as well as up to four constraints on its radiative forcing (though only one on the RF of each of the (H)CFCs if we assume they are acting in the weak limit as discussed in Section 6.2.2).

I could reduce the number of constraints by fixing some of the variables as constant, for example in Equation 6.4, the radiative forcing equation for methane and N_2O, I could possibly set one or two of the terms to constant. This was done in the carbon cycle of the box-diffusion model to avoid there being too many free parameters. I would still probably

need one, two or possibly even three free parameters in Equation 6.4 in order to represent the full range of uncertainty in the forcing due to these species.

I could potentially reduce the number of observational constraints required by taking results from model output, as we did when using the C^4MIP output to constrain the carbon cycle of the box-diffusion model. However it is not obvious which results to take.

In additions to the question of which constraints to use, coding this up in the box-diffusion model would probably take months of work (Kahneman & Tversky 1977). This is time that I chose instead to use developing my work on SLCFs, which is presented in Chapter 7.

Having introduced the non-CO_2 components that I have added to the box-diffusion model, I will now perform two simple tests of the non–CO_2 components of the box-diffusion model.

6.5 Testing our treatment of non-CO_2 greenhouse gases

6.5.1 Comparison using metrics

Using the metrics introduced in Section 3.5, we can compare the box-diffusion model's response to a pulse emission against the response used in the IPCC AR5. This is done in Table 6.3. The white rows contain data taken from the IPCC AR5, calculated using the Bern carbon cycle model (Joos *et al.* 1996; Myhre & Shindell 2013), and the impulse response climate model from Boucher & Reddy (2008). The grey rows contain data calculated by the box-diffusion model. The data for CO_2 is the same as that presented in Chapter 3. The data for the non-CO_2 GHGs is calculated using the simple models introduced in this chapter.

In calculating these metrics, I and the IPCC have only considered the feedbacks of CO_2 on itself. I have not considered CO_2 feedbacks on other gases. If CO_2 feedbacks on methane and N_2O are included, their values are likely to increase GWPs by a factor of 20%, and the GTP-100 of methane by a factor of approximately 100% (Gillett & Matthews

Trace gas	**Lifetime** (years)	**Radiative efficiency** ($Wm^{-2}ppb^{-1}$)	**AGWP** *20-year* (Wm^{-2} yr kg^{-1})	**GWP** *20-year*	**AGWP** *100-year* (Wm^{-2} yr kg^{-1})	**GWP** *100-year*	**AGTP** *20-year* (K kg^{-1})	**GTP** *20-year*	**AGTP** *100-year* (K kg^{-1})	**GTP** *100-year*
CO_2	N/A	1.4×10^{-5}	2.5×10^{-14}	*1*	9.2×10^{-14}	*1*	6.9×10^{-16}	*1*	5.4×10^{-16}	*1*
	N/A	1.4×10^{-5}	1.8×10^{-14}	*1*	8.6×10^{-14}	*1*	4.6×10^{-16}	*1*	6.2×10^{-16}	*1*
CH_4	12	3.6×10^{-4}	2.0×10^{-12}	*84*	2.5×10^{-12}	*28*	4.5×10^{-14}	*65*	2.3×10^{-15}	*4*
	9	3.7×10^{-4}	1.1×10^{-12}	*60*	1.2×10^{-12}	*14*	1.7×10^{-14}	*37*	6.3×10^{-15}	*1*
N_2O	121	3.0×10^{-3}	7.2×10^{-12}	*286*	2.6×10^{-12}	*288*	2.1×10^{-13}	*301*	1.4×10^{-13}	*254*
	121	3.2×10^{-3}	7.6×10^{-12}	*430*	2.8×10^{-11}	*320*	2.1×10^{-13}	*450*	1.4×10^{-13}	*220*
CFC-11	45	0.26	1.7×10^{-10}	*6957*	4.3×10^{-10}	*4702*	4.7×10^{-12}	*6942*	1.3×10^{-12}	*2353*
	45	0.26	1.8×10^{-10}	*10000*	4.3×10^{-10}	*5200*	4.4×10^{-12}	*10000*	1.2×10^{-12}	*2000*
CFC-12	100	0.32	2.7×10^{-10}	*10814*	9.4×10^{-10}	*10257*	7.7×10^{-12}	*11278*	4.6×10^{-12}	*8458*
	100	0.32	2.8×10^{-10}	*16000*	9.5×10^{-10}	*12000*	7.3×10^{-12}	*17000*	4.3×10^{-12}	*7000*
CFC-113	85	0.3	1.6×10^{-10}	*6432*	5.3×10^{-10}	*5771*	4.6×10^{-12}	*6665*	2.4×10^{-12}	*4430*
	85	0.3	1.6×10^{-10}	*9600*	5.3×10^{-10}	*6500*	4.3×10^{-12}	*10000*	2.3×10^{-12}	*3800*
CCl_4	26	0.18	8.9×10^{-11}	*3588*	1.6×10^{-10}	*1780*	2.3×10^{-12}	*3373*	2.7×10^{-13}	*493*
	26	0.18	9.1×10^{-11}	*5300*	1.6×10^{-10}	*2000*	2.1×10^{-12}	*4800*	2.3×10^{-13}	*390*
HCFC-22	12	0.2	1.2×10^{-10}	*5049*	1.5×10^{-10}	*1687*	2.7×10^{-12}	*4009*	1.4×10^{-13}	*250*
	12	0.2	1.3×10^{-10}	*7500*	1.6×10^{-10}	*1900*	2.3×10^{-12}	*5300*	1.0×10^{-13}	*160*

Table 6.3: Metrics for the different gases incorporated into the box-diffusion model. White rows contain numbers from the IPCC Fifth Assessment Report (Myhre & Shindell 2013), given to the number of significant figures presented in the report. Grey rows contain numbers obtained from the box-diffusion model, and are given to two significant figures. Numbers are calculated relative to present day concentrations. CO_2 carbon cycle feedbacks are included in the calculation of metrics for CO_2, but not for the other GHGs. CO_2 lifetime is not given as it does not have a single lifetime as discussed in Chapter 2.

2010; Myhre & Shindell 2013). Other effects not included in these estimates include the ozone-CO_2 effect, which would increase GTPs of methane by a factor of two, and oxidant-sulphate interactions which may also increase the GWP-100 of methane by up to 30% (Myhre & Shindell 2013).

Differences between the IPCC numbers and the box-diffusion numbers occur in a few places, as expected. The CO_2 impulse responses are quite different, as described in Chapter 3. The behaviour of methane is different because the IPCC uses a 12-year lifetime rather than the 9-year lifetime that I chose to use following Holmes *et al.* (2013). Because of the relatively large differences in the behaviour of CO_2, the GWPs and GTPs of the non-CO_2 GHGs are also different, despite the fact that their AGWPs and AGTPs are quite similar. There are no unexpected differences between the IPCC metrics and those calculated by the box-diffusion model, so I will go on to test the performance of the aggregate of the box-diffusion's non-CO_2 components against the CMIP5 ensemble.

6.5.2 Comparison with CMIP5

In Figure 6.7 I compare the warming simulated by the box-diffusion model against that simulated by CMIP5. This plot is very similar to Figure 3.3. The dotted lines in this plot are the same as the dashed lines in Figure 3.3a. The dashed lines in this plot are the same as the dashed lines in Figure 3.3b. The differences between these two lines are discussed in Section 3.3. The line that has been added to this plot is the very thick solid line, representing the temperature projection of the RCPs driven by the eight GHG emissions pathways simulated in the box-diffusion model. This suggests that there is a very slight reduction in the warming in our model due to non-CO_2 GHGs compared with MAGICC which was used to create the unofficial RCP emissions scenarios (Meinshausen *et al.* 2011a). This is likely to be due to the difference in the treatment of methane between MAGICC and the box-diffusion model. The difference is rather small though, and the results are still quite close to the multi-model mean of the CMIP5 ensemble, and well within the spread of the CMIP5 ensemble.

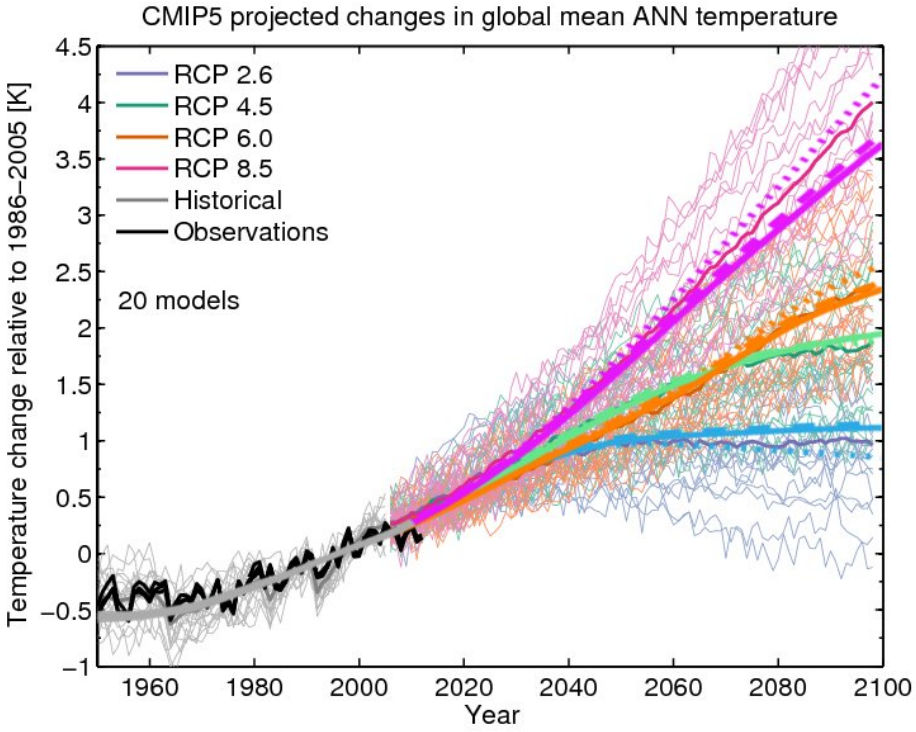


Figure 6.7: Comparing the box-diffusion model with CMIP5 model output. Thin lines are individual CMIP5 models (Taylor *et al.* 2012). Thicker wiggly lines are the multi-model means across CMIP5. Dotted thick lines correspond to the box-diffusion model forced with the unofficial RCP forcings (Meinshausen *et al.* 2011c). Dashed lines correspond to the box-diffusion model forced with CO_2 emissions, and unofficial RCP forcings for the non-CO_2 GHGs. The thick solid lines corresponds to the box-diffusion model forced by unofficial RCP emissions for both CO_2 and non-CO_2 GHGs (Meinshausen *et al.* 2011c). Colours correspond to the RCP emissions scenarios introduced in Section 2.5. Grey lines are historic model output. Black lines are historical observations (Lawrimore *et al.* 2011; Hansen *et al.* 2010; Jones *et al.* 2012). CMIP5 models use the official RCP concentrations as inputs (Meinshausen *et al.* 2011c; van Vuuren *et al.* 2011). Source: Adapted from an unpublished plot by Ed Hawkins (Met Office).

These results, combined with the others presented in this chapter, suggest that we can now go on to use the box-diffusion model to consider the differences between SLCFs and LLCFs in Chapter 7.

Chapter 7

Short-lived climate forcers

7.1 Short- and long-lived climate forcers

In chapters 4 and 5 I established that it is the cumulative emission of CO_2 that determines peak CO_2-induced warming. However, as discussed in Chapter 6 CO_2 is not the only GHG that impacts on peak warming, as illustrated in Figure 2.7. Emissions of both CO_2 and other GHGs are contributing to global warming (Forster *et al.* 2007). It would be useful to establish what property of the emissions of these other GHGs controls their contribution to peak warming. In a paper that I co-authored with Smith *et al.* (2012) we did just that.

In Figure 7.1 we plotted two correlations: on the y-axis the correlation between the maximum emission rate of a gas and the peak resultant temperature as illustrated for CO_2 in Figure 4.1f, and on the x-axis the correlation between cumulative emissions of a gas and its peak resultant temperature as illustrated for CO_2 in Figure 4.1a. By conducting an analysis not dissimilar to that outlined in Chapter 4 we found that the gases naturally divide themselves into two categories: gases for which resultant peak warming correlates better with the peak emissions rate ($r^2(\Delta T_{\max}, \sum E) < r^2(\Delta T_{\max}, E_{\max})$), and gases for which resultant peak warming correlates better with cumulative emissions ($r^2(\Delta T_{\max}, \sum E) > r^2(\Delta T_{\max}, E_{\max})$). Let us define these two groups as short-lived climate forcers (SLCFs) and long-lived climate forcers (LLCFs) respectively. The SLCFs are in the upper and left half of Figure 7.1 and the LLCFs are in the lower and right half. As shown in the Figure, SLCFs generally have lifetimes of less than ~30 years, while LLCFs

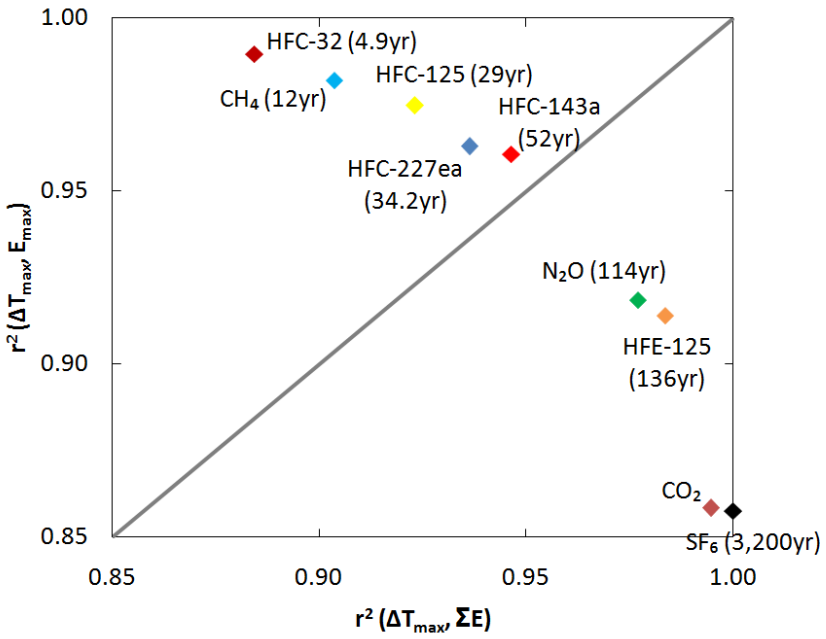


Figure 7.1: Correlation between peak temperature rise and cumulative emissions $(r2(\Delta T_{\max}, \sum E))$ versus correlation between peak temperature rise and maximum emission rate $(r^2(\Delta T_{\max}, E_{\max}))$ for a subset of GHGs. Line of equal correlations shown in grey. Source: Smith *et al.* (2012)

have lifetimes of at least ~100 years. This timescale is probably set by a combination of the thermal response timescales of the climate system and the timescales of the scenarios used to generate this figure. Only hydrofluorocarbon-143a (HFC-143a; lifetime 52 years) and HFC-227ea (lifetime 34.2 years) lie between these two thresholds, and they account for less than 1% of the total RF at any point in any of the RCPs (Meinshausen *et al.* 2011c).

The thesis of Smith *et al.* (2012) is that the warming at time t can be estimated by calculating warming caused by the cumulative emissions of LLCFs up to time t and adding it to the warming caused by the emissions of SLCFs at time t.

Having defined SLCFs and LLCFs I can now discuss the proposals to reduce emissions of SLCFs .

7.2 Action on short-lived climate forcers

Recent years have seen continued growth in emissions of CO_2 (Canadell *et al.* 2007; Raupach *et al.* 2007, 2008; Le Quéré *et al.* 2009; Friedlingstein *et al.* 2010; Peters *et al.* 2011a, 2012), as well as increased interest in other, shorter-lived climate forcing agents such as black carbon, tropospheric ozone and methane (Molina *et al.* 2009; Ramanathan & Xu 2010; UNEP & WMO 2011; Shindell *et al.* 2012; Wallack & Ramanathan 2009). A new initiative to tackle SLCF emissions has been created by the Climate and Clean Air Coalition to reduce Short-Lived Climate Pollutants, including all of the G8 nations. The emissions reductions would be made via a number of interventions outlined in UNEP & WMO (2011) including the recovery of methane from coal, oil and gas extraction and transport, methane capture in waste management, use of clean-burning stoves for residential cooking, diesel particulate filters for vehicles and the banning of field burning of agricultural waste. There are substantial co-benefits, not related to global climate, of SLCF mitigation, including benefits to agriculture and human health (Shindell *et al.* 2012; UNEP & WMO 2011; Wallack & Ramanathan 2009). Our focus here is purely on the potential impact of SLCF mitigation on global climate.

Two of the most common ways of understanding the aims of mitigation policy have been climate stabilisation (represented by limiting either greenhouse gas concentrations or global mean temperature) and limiting the rate of climate change to a level at which societies and ecosystems might adapt (United Nations 1992; UNFCCC 2011; Peck & Teisberg 1994; Parry *et al.* 1996; Root *et al.* 2003; Kallbekken *et al.* 2009). Much of the discussion of the climate benefits of near-term reductions in short-lived climate pollutants (SLCFs) has been around their ability to reduce the near-term rate and magnitude of warming, and their ability to delay the crossing of temperature thresholds (Molina *et al.* 2009; Ramanathan & Xu 2010; UNEP & WMO 2011; Shindell *et al.* 2012; Wallack & Ramanathan 2009; Rypdal *et al.* 2009).

The main argument presented in this chapter is that reducing SLCFs will not help to meet climate stabilisation targets unless emissions of long-lived climate pollutants (LLCFs) such as carbon dioxide (CO_2) are falling simultaneously. The secondary argument I present is that reductions in SLCFs will not help reduce the maximum rate of warming unless they are timed to coincide with the period of peak LLCF emissions. My final argument is an alternative presentation of our primary point: SLCF reductions cannot be used to buy time to delay reductions in long-lived climate forcings (LLCFs) such as CO_2. Before I discuss these points, I will outline how the SLCF reductions are modelled in this work.

7.3 Methods

The modelling presented in this chapter is carried out using the box-diffusion model used throughout this thesis. This work consists of carrying out a cut in emissions of SLCFs, as shown in Figure 7.2.

I use the RCPs as baseline scenarios, using the harmonised emissions and forcing data from Meinshausen *et al.* (2011c). CO_2 emissions in RCP-4.5 were smoothed between 2070 and 2110 before being used in our analysis, as I could not justify CO_2 emissions rising in absolute terms between 2080 and 2100 after rapid reductions in the previous decades (Meinshausen *et al.* 2011c). Based on the proposals put forward by UNEP & WMO (2011) I reduce black carbon forcings and methane emissions linearly over the course of two decades to 78% and 25% respectively relative to the RCP baseline scenarios. I keep the black carbon and methane emissions reduced to this fraction of the baseline scenario indefinitely after the cuts have been made. I do not alter sulphur dioxide (SO_2), organic carbon or (H)CFC emissions from those in the RCP baseline scenarios, since these are not targeted in the UNEP/WMO SLCF mitigation proposals (UNEP & WMO 2011). When I discuss (H)CFCs throughout this thesis I include carbon tetrachloride (CCl_4), as well as the other HCFCs and CFCs included in Table 6.1.

The impact of maximal mitigation of tropospheric ozone (UNEP & WMO 2011) is represented in the same fashion as black carbon, namely at the forcing level. For tropospheric ozone I halve the radiative forcing, this being the ratio of 2010 tropospheric ozone forcing to long-term future ozone forcing in RCP-3PD. Results presented in this chapter are not sensitive to the magnitude of the tropospheric ozone reduction.

The overall reduction in SLCF forcing is 1.1 Wm^{-2} if the cuts were made completely in 2010. The reduction in radiative forcing relative to the baseline scenario decreases over time in absolute terms, because the SLCF measures are modelled as reductions relative to the RCP baseline scenarios, and emissions of SLCFs in the baseline scenario decrease over time. By 2050 the reduction in radiative forcing due to SLCF measures relative to the RCP baseline scenarios is 0.7 Wm^{-2} in RCP-3PD, 0.9 Wm^{-2} in RCP-4.5, 1.0 Wm^{-2} in RCP-6.0 and 1.0 Wm^{-2} in RCP-8.5. These reductions are larger than those presented by Ramanathan & Xu (2010) and Shindell *et al.* (2012) because of recent evidence that the current radiative forcing due to black carbon is larger than previously thought (Bond *et al.* 2013), as discussed in Chapter 6. The results presented in this chapter are not sensitive to a smaller BC forcing, nor a longer methane lifetime, as discussed in Chapter 6.

In order to avoid showing only temperature response curves calculated with best guess model parameters, I have also included low- and high-response plots for each figure that I have made in this chapter. The method of selecting low- and high-response parameters is outlined in Section 2.6.2.

7.4 Delaying SLCF reductions

Figures 7.2 and 7.3 show the SLCF reductions proposed by UNEP & WMO (2011), and their impact. For each baseline, and for each set of response parameters, the plots show sustained SLCF reductions reduce peak warming relative to no reductions, but they also show that delaying the emissions reductions to a future decade would not significantly impact peak warming.

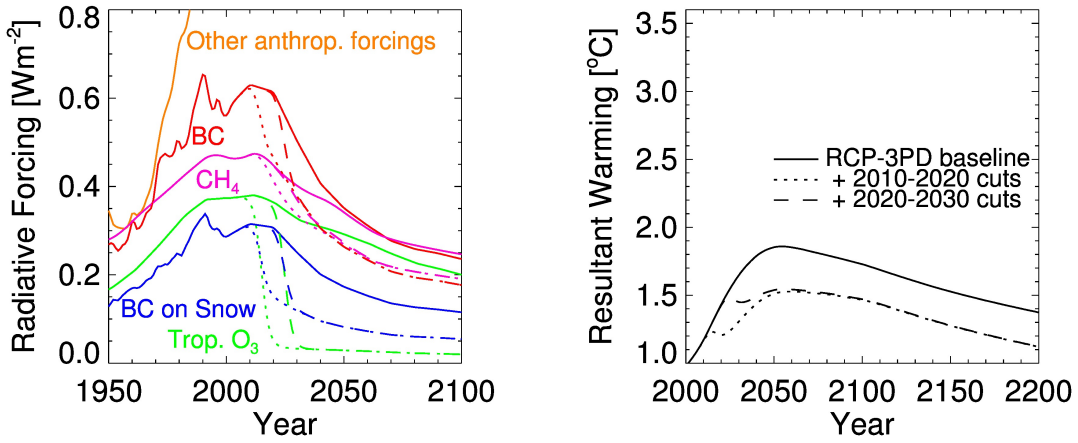
I show only the two of the four RCP baseline scenarios with the lowest total cumulative carbon emissions, as these are the two with the lowest peak warmings. The same point can be made about RCP-6 and RCP-8.5 as shown in Figures 7.7 and 7.8.

Much of the remainder of the analysis in this chapter is based on this simple observation that delaying SLCF reductions to a future decade will not impact on peak warming under these scenarios.

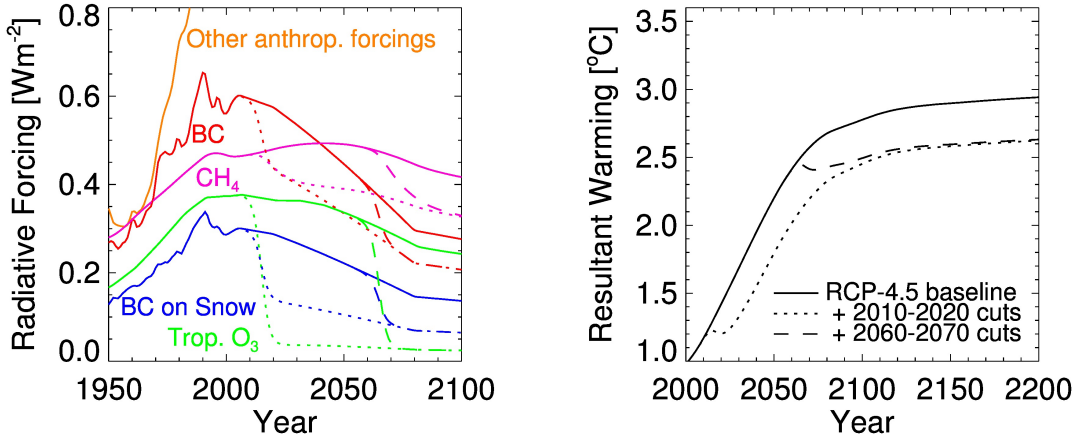
7.5 How much do short-lived climate forcers matter for peak warming?

Reducing emissions of non-CO_2 GHGs such as methane and nitrous oxide will help mitigate climate change via global radiative forcing or climate forcing, which controls the magnitude of global warming. One of the papers that I studied during my DPhil was Montzka *et al.* (2011). The paper appears to mistakenly imply that reductions in short-lived GHGs emitted today will reduce peak radiative forcing, which even on the most optimistic mitigation scenarios, will occur in several decades' time. The authors must have been referring to sustained emissions reductions of SLCFs reducing peak radiative forcing. Short-lived GHGs remain in the atmosphere for only a few years (Forster *et al.* 2007), so emissions of these gases over the coming decade cannot substantially affect either peak radiative forcing, or consequently the magnitude of peak warming, under scenarios in which temperatures peak around at least 2°C above pre-industrial. Current emissions of short-lived gases have even less impact on peak temperatures under scenarios in which temperatures peak later and higher. Rather, it is emissions of short-lived GHGs near the time of peak radiative forcing that significantly affect that peak in radiative forcing and subsequently peak warming (Shine *et al.* 2007). Short-lived GHGs emitted today can indirectly affect peak warming through carbon-cycle feedbacks, but this is a secondary effect, as shown in Figure 7.4.

Figure 7.4 shows the contribution to peak warming from various forcing agents as a

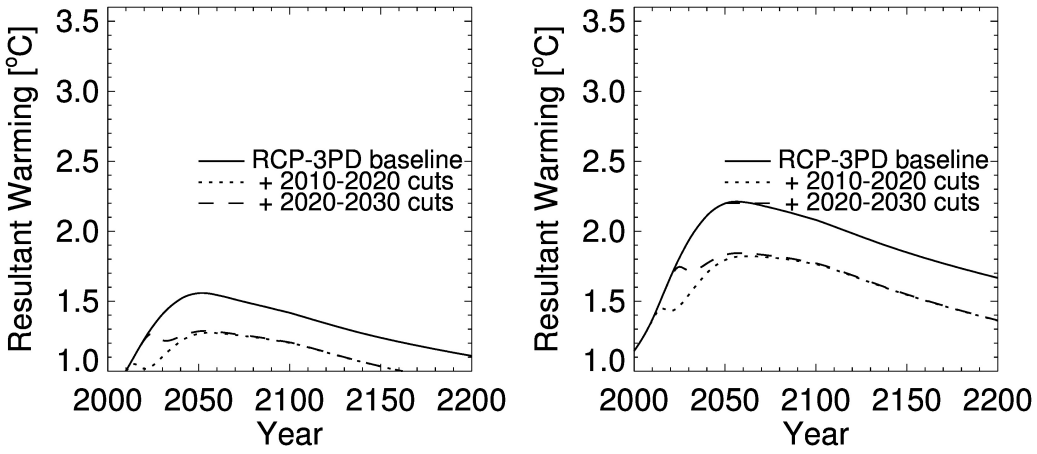


(a) RCP-3PD used as the baseline scenario with best-guess carbon cycle and climate parameters

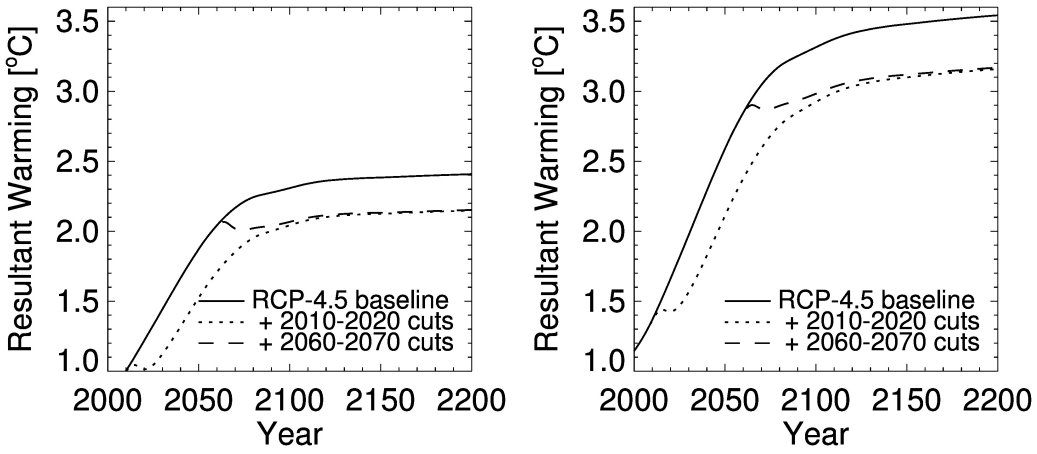


(b) RCP-4.5 used as the baseline scenario with best-guess carbon cycle and climate parameters

Figure 7.2: Temperature response to different scenarios. The solid lines represent the baseline scenarios. The dotted lines have cuts in short-lived species similar to those proposed by Shindell *et al.* (2012), while the dashed line has the same fractional cuts postponed by (a) 10 years and (b) 50 years. Other anthropogenic forcings remain the same across all scenarios in each figure. The uncertainty on this figure is shown in Figure 7.3.

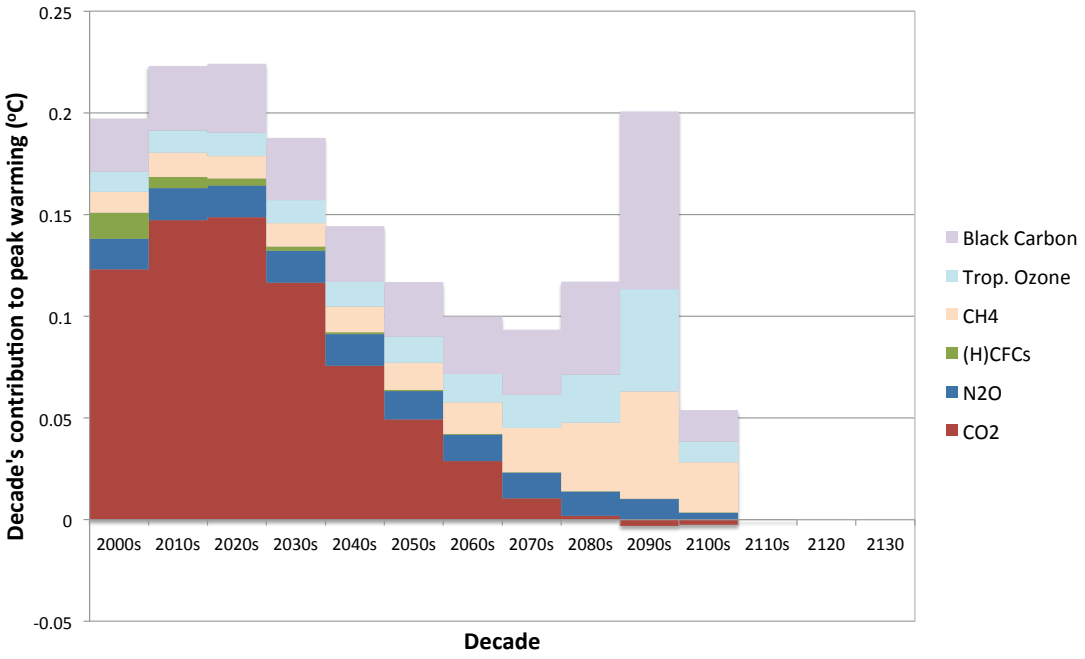


(a) RCP-3PD with low response parameters **(b)** RCP-3PD with warm model parameters

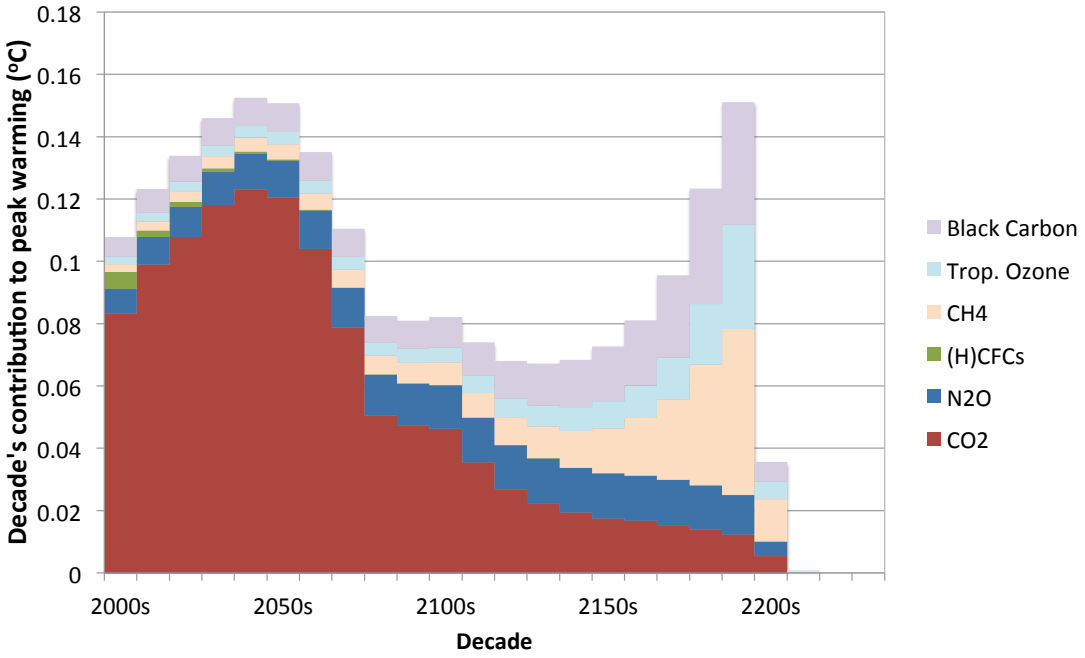


(c) RCP-4.5 with low response parameters **(d)** RCP-4.5 with warm model parameters

Figure 7.3: Uncertainty on Figure 7.2. Panels (a) and (b) are as the right hand panels in Figure 7.2a, but with low and high response model parameters respectively. Panels (c) and (d) are the same for Figure 7.2b. Together, these two model parameter setting span the 5-95% confidence interval on the response in the box-diffusion model.

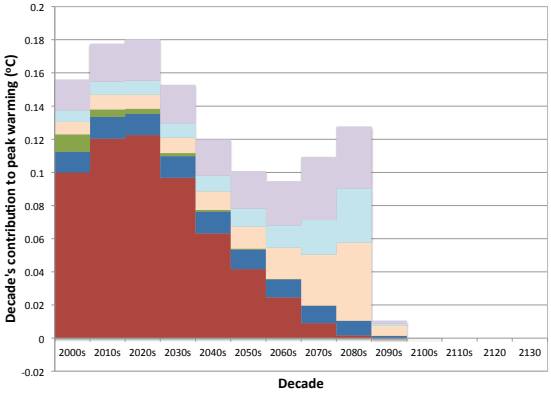


(a) RCP-3PD used as the baseline scenario with best-guess model parameters

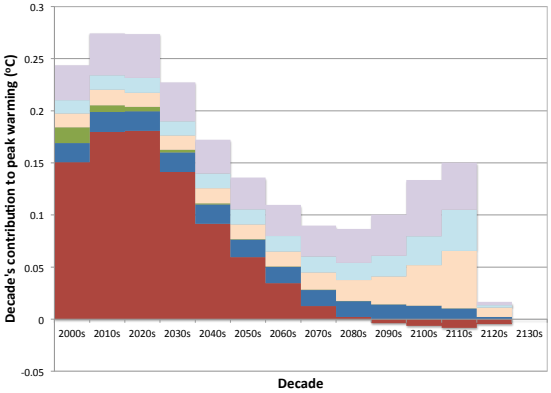


(b) RCP-4.5 used as the baseline scenario with best-guess model parameters

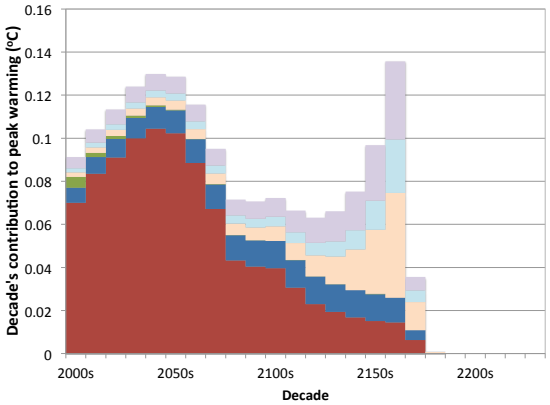
Figure 7.4: Contribution of decades' emissions to peak warming under (a) RCP-3PD and (b) RCP-4.5. To create these plots I have removed the emissions of a given greenhouse gas in a given decade from a model run and calculated the resulting impact of these emissions on peak global warming. These plots us the best-guess model parameters. Uncertainty on this plot is shown in Figure 7.5.



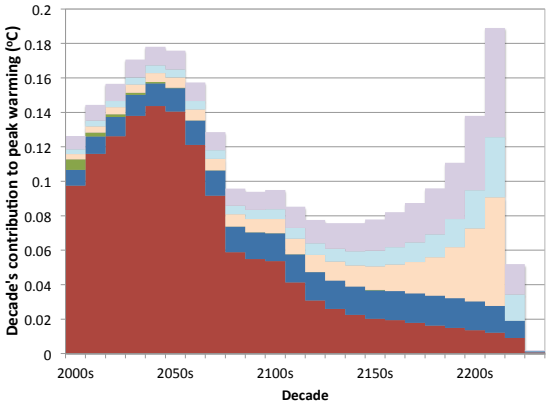
(a) RCP-3PD with low response parameters



(b) RCP-3PD with low response parameters



(c) RCP-4.5 with low response parameters



(d) RCP-4.5 with high response parameters

Figure 7.5: Uncertainty on Figure 7.4. Panels (a) and (b) are as Figure 7.4a, but with low and high response parameters respectively. Panels (c) and (d) are the same for Figure 7.4b. These two model parameter setting span the 5-95% confidence interval on the response in the box-diffusion model. Note that the scales in the plots are different, in order to make it easier to compare the overall shapes of the different responses.

function of the decade of emission, calculated simply by omitting the relevant decade's emissions of a particular forcing agent and computing the resulting change in peak warming. While CO_2 emissions are dominant initially, their importance declines as CO_2 emissions decline in these two scenarios, and the importance of short-lived agents increases significantly near the time of peak warming. I use two emissions scenarios, RCP-3PD and RCP-4.5, which I find cause global mean temperatures to peak at approximately 2°C and 3°C respectively with the most likely values of model parameters. The results for low- and high-response parameters are shown in Figure 7.5. Though the details change between scenarios, the qualitative result remains the same: in terms of current emissions, CO_2 has a far greater impact on peak radiative forcing and peak warming than short-lived GHGs for any plausible emissions scenario.

Short-lived forcing agents emitted today raise global temperatures temporarily, causing increased emission of CO_2 from the biosphere via carbon-cycle feedback mechanisms (Gillett & Matthews 2010). Through this carbon-cycle feedback, today's emissions of short-lived GHGs can impact peak forcing, and in turn peak warming. However, as shown in Figure 7.4 this is a secondary effect that has a much smaller impact in any given decade than the influence of emissions of short-lived forcing agents near the time of peak forcing.

7.5.1 Reproducing results

Because Figure 7.4 was produced in the very simple box-diffusion model I wanted to see if it could be reproduced by MAGICC. Steve Smith kindly re-ran the experiment using MAGICC as part of the paper that we wrote together (Smith *et al.* 2012). The results, plotted in Figure 7.6, are similar to the equivalent plot in Figure 7.4a, though there are some key differences. MAGICC takes CO_2 out of the atmosphere faster than the box-diffusion does as explained in Chapter 2, and so temperatures peak sooner, which is responsible for most of the difference. The forcing due to black carbon in the Smith *et al.* (2012) plot is approximately three times too small compared to the estimates by Bond *et al.* (2013), and compared to the estimates used in Figure 7.4.

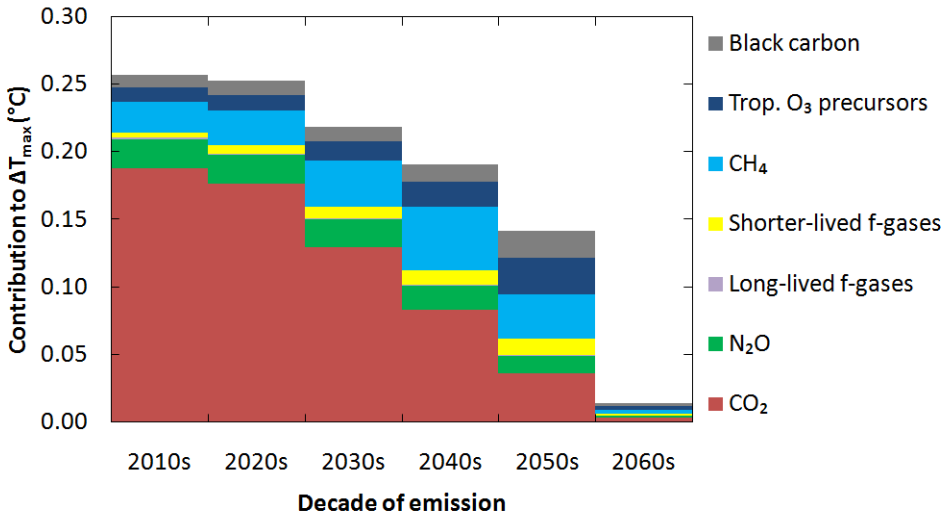


Figure 7.6: Contributions to peak temperature in RCP-3PD from emissions of GHGs, black carbon and tropospheric ozone precursors in each decade. Tropospheric ozone precursors denote emissions of carbon monoxide, oxides of nitrogen and volatile organic compounds and includes the indirect effect on methane through OH concentrations. F-gases denote halogen-containing compounds. CH_4 contribution includes indirect effects on tropospheric ozone and stratospheric water vapour; black-carbon contribution includes the effect on snow albedo. For this model configuration and all anthropogenic forcings in the RCP-3PD scenario, peak warming occurs in 2062. Source: Smith *et al.* (2012)

Because MAGICC takes carbon out of the atmosphere faster, CO_2 emissions are still falling around the time of temperature peak, whereas CO_2 emissions are close to zero and in fact negative when temperatures peak in the box-diffusion model. This means that the increase in the response due to the SLCFs black carbon, tropospheric ozone and methane is more easily visible in Figure 7.4 than it is in Figure 7.6 where it is disguised by the falling contribution from falling CO_2 emissions. The magnitude of change in the response from the SLCFs is similar, with the response from black carbon and tropospheric ozone increasing by a factor of approximately three over the five decades in the run up to peak warming. If the magnitude of the response in Figure 7.4 is larger, it is likely to be due to the fact that it is implemented at the forcing level, whereas changes in Figure 7.6 are implemented at the emissions level, giving them a slightly longer effective lifetime.

7.6 The role of short-lived climate forcers in meeting temperature goals

Throughout this thesis I have been considering emission targets to avoid dangerous climate change. One of the most commonly used definitions of dangerous climate change is a 2°C rise in global average temperatures relative to pre-industrial levels (UNFCCC 2010). The contents of this section have been accepted as part of a perspectives article for *Nature Climate Change*, and hence this section has a more policy-oriented focus and tone than the remainder of this thesis.

Early action on SLCF emissions has been criticised for potentially delaying action on CO_2 either by diverting resources or on the assumption that action only occurs after temperatures cross some threshold (Myhre *et al.* 2011; Asrar & Hurrell 2012). There have been counter-arguments that claim that effective action on SLCFs could help build political credibility and thus mobilise additional resources for action on other pollutants, notably CO_2 (UNEP & WMO 2011; Shindell *et al.* 2012; Wallack & Ramanathan 2009; Victor *et al.* 2012). I do not propose to enter this debate here, but instead address a simpler question: do current SLCF emissions have any impact at all on the risk of crossing potentially dangerous temperature thresholds? I show that, in the absence of a near-term reversal of CO_2 emission trends, they do not. Early emphasis on SLCF reduction without simultaneous reductions in CO_2 emissions will make it difficult to avoid crossing a temperature threshold such as the two-degree target. My arguments apply generically to other temperature targets as well, such as the 1.5°C target recently discussed by the UNFCCC (UNFCCC 2011).

7.6.1 Buying time for adaption but not mitigation

The impact of the proposed approximately 25% and 78% emissions reductions in methane and black carbon emissions (UNEP & WMO 2011) respectively over a two-decade period on most likely global temperatures for a range of start dates under four standard RCP scenarios (van Vuuren *et al.* 2011; Meinshausen *et al.* 2011c) are shown in Figures 7.7 and

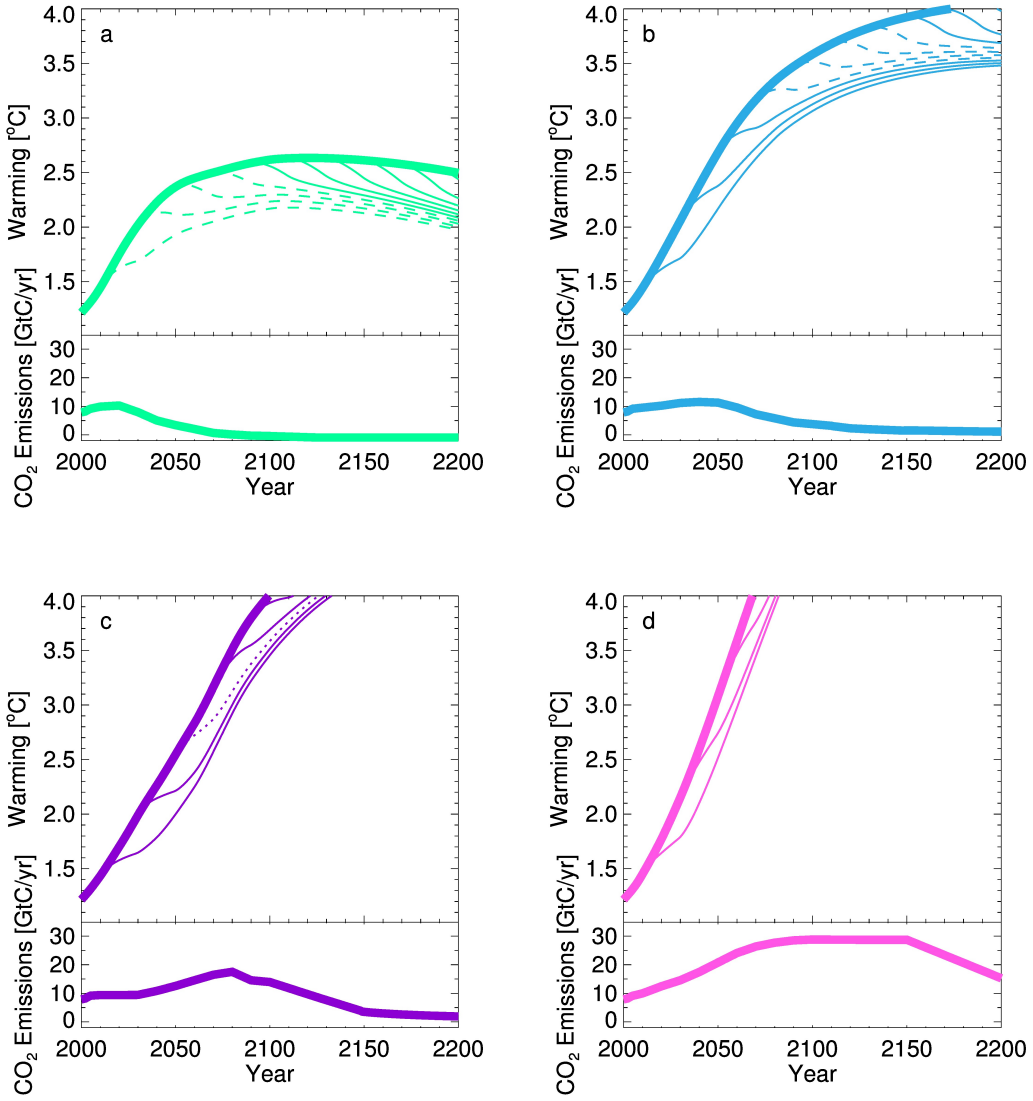
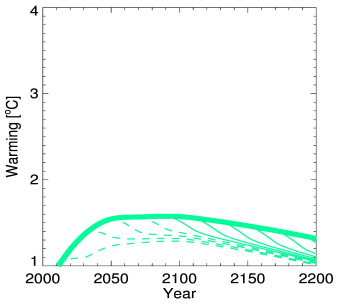
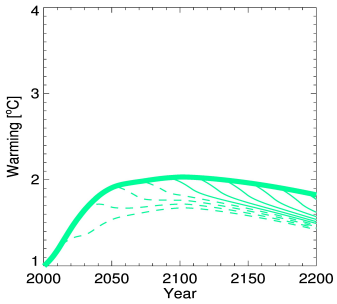


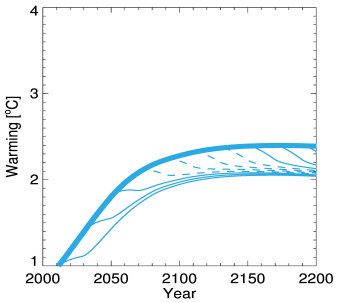
Figure 7.7: Global warming from implementing SLCF measures proposed by UNEP & WMO (2011) in different decades and for four different baseline emission scenarios. Thick lines indicate baseline scenarios, and the branching lines illustrate the impact of implementing SLCF cuts in every second decade from the 2010s onwards. Dashed lines correspond to SLCF cuts that have more than 0.06°C impact on peak warming relative to delaying the SLCF measures by two decades, while thin solid lines correspond to cases where SLCF cuts have less than 0.06°C impact on this measure. Dotted scenario in subfigure c corresponds to gradual SLCF reductions to minimise peak warming rates as discussed in the text. Each line plotted here corresponds to a coloured point plotted in Figure 7.9. SLCF cuts make little difference to peak warming relative to making the same cuts two decades later, except in the limited number of cases plotted as dashed lines which coincide with periods of rapid SLCF emission reductions (lower panels). This plot uses the best-guess model parameters in the box-diffusion model, and the uncertainty on this plot is shown in Figure 7.8.



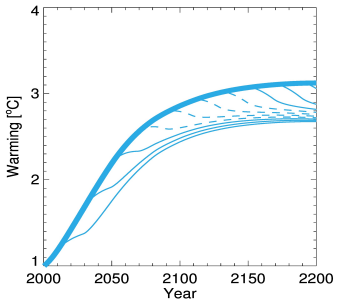
(a) RCP-3PD with low response parameters



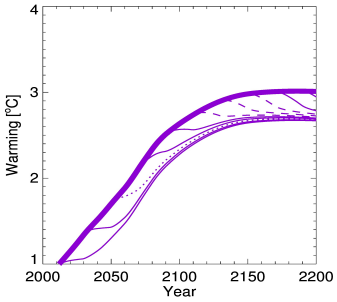
(b) RCP-3PD with high response parameters



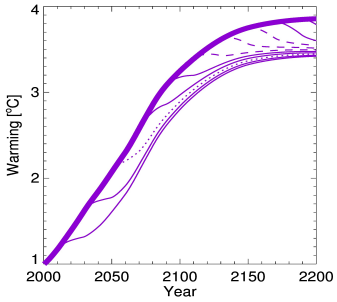
(c) RCP-4.5 with low response parameters



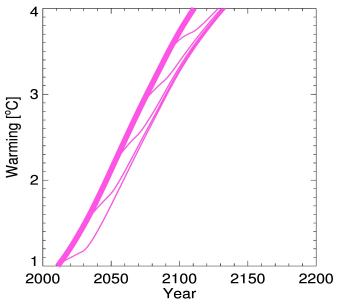
(d) RCP-4.5 with high response parameters



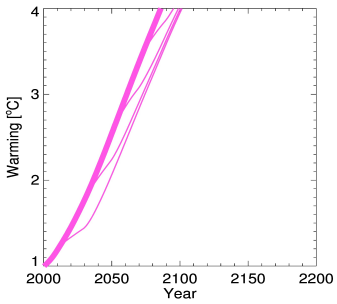
(e) RCP-6 with low response parameters



(f) RCP-6 with high response parameters



(g) RCP-8.5 with low response parameters



(h) RCP-8.5 with high response parameters

Figure 7.8: Uncertainty on Figure 7.7. Panels (a), (c), (e) and (g) are the same as Figure 7.7 but use low response parameters. Panels (b), (d), (f) and (h) are the same, but using warm model parameters. These two model parameter setting span the 5-95% confidence interval on the response in the box-diffusion model.

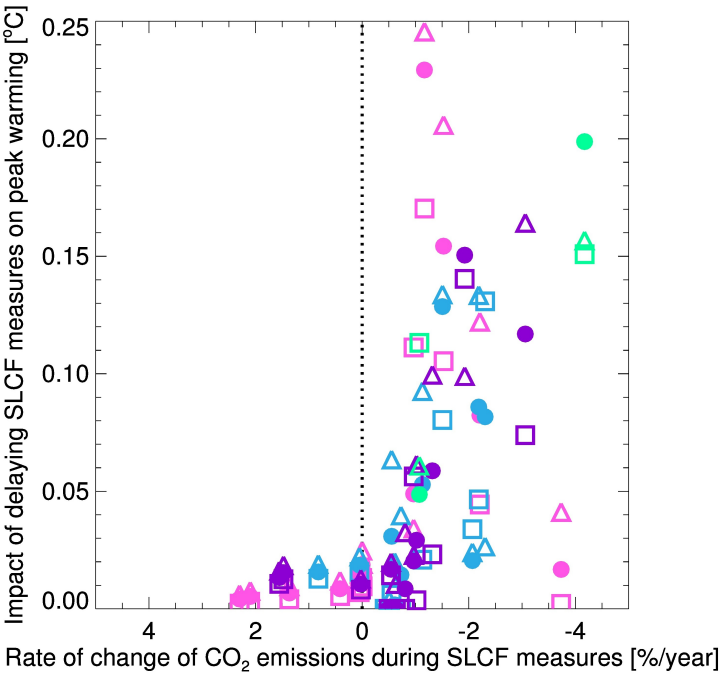


Figure 7.9: Impact on peak warming of early implementation of SLCF measures versus simultaneous rate of reduction of CO_2 emissions. Vertical axis shows the most likely increase in peak warming resulting from delaying SLCF measures by two decades, as calculated from the warming projections in Figure 7.7, while the horizontal axis shows the average rate of change of CO_2 emissions over the two decades in which the SLCF measures are first implemented, calculated from the baseline CO_2 emissions scenarios in Figure 7.7. Colours correspond to those used in Figure 7.7. The three symbols correspond to three different parameter combinations spanning the IPCC likely range. Triangles correspond to high response parameters, and squares correspond to low response parameters. Filled circles correspond to the best guess parameter set, as discussed in Section 2.6.2. Decades with negative net CO_2 emissions are not plotted.

7.8. While the potential impact of SLCF reduction measures on lowering peak temperatures is clear, the case for early action is less obvious. Only under specific conditions, illustrated by the dashed lines in Figures 7.7 and 7.8, do these measures make more than a few hundredths of a degree difference to peak temperatures relative to the impact of taking identical measures two decades later. Under the ~2°C scenario (in green) SLCF reductions will begin to matter for peak warming over the next couple of decades as CO_2 emissions start falling rapidly. Under the other scenarios, however, it makes almost no difference whether the SLCF cuts are implemented now or deferred until after 2060.

This matters, because early action on SLCFs has been widely reported as an option for "buying time" (Ramanathan & Victor 2010; Morton 2011; Biello 2012) for implementing reductions in LLCFs. Figure 7.9 shows the impact on peak warming of implementing these SLCF measures in any given decade, relative to deferring them by two decades (their hypothetical implementation time). This impact has been plotted as a function of the rate of reduction of CO_2 emissions over the period in which SLCF measures are implemented for low (squares), best guess (circles) and high (triangles) response parameters as described in Section 2.6.2. A similar plot to Figure 7.9 is obtained (not shown) if we plot the impact of delaying SLCF measures against the rate of reduction of total LLCF emissions. Proposed reductions in SLCFs only have a significant impact on peak warming if they are implemented over a period when CO_2 (or total LLCF) emissions are falling.

Even under the 2°C stabilisation scenario, the combined impact of methane and black carbon emissions over the decade 2010 to 2020 is to increase the most likely peak warming by less than a few hundredths of a degree, as illustrated in Figures 7.4 and 7.6. In contrast, LLCF emissions over this period will contribute of order ten times that amount. The reason is that by the time temperatures peak, the impact of current SLCF emissions will have largely dissipated (precisely because they are short-lived), while the impact of LLCFs emitted now persists for centuries (Shine *et al.* 2005, 2007; Solomon *et al.* 2009; Matthews & Caldeira 2008; Matthews & Weaver 2010).

Under the 2°C scenario, which requires rapid reductions in CO_2 emissions over the next 20 years (Anderson 2008; Anderson *et al.* 2008; Anderson & Bows 2011; Vuuren *et al.* 2011; Rogelj *et al.* 2011b, 2012b), the impact of SLCF emissions rises rapidly over the next few decades. Since it takes time to negotiate treaties and put the technology in place to make deep cuts in SLCFs possible (Friedlingstein *et al.* 2011; Peters *et al.* 2012), immediate action on SLCFs may be justified as a means of limiting peak warming to meet the 2°C goal, but only if we are confident in our ability and willingness to reduce CO_2 emissions rapidly over the same period. Under scenarios in which temperatures peak at 3°C or 4°C, however, SLCF emissions only begin to have a substantial impact on peak warming much later this century. Hence reductions in SLCFs would prove effective in the next few decades if we implement the required LLCF measures to meet the 2°C goal, but they will not matter much for several decades if we do not.

A further argument for early reductions in SLCFs is that they help reduce the rate of warming (Wallack & Ramanathan 2009; Ramanathan & Xu 2010; UNEP & WMO 2011; Shindell *et al.* 2012). Here we argue that while SLCF reductions would reduce near-term rates of warming, they will not necessarily reduce the maximum rate of warming. If peak warming is to be limited to 2°C, the relatively small remaining allowable emissions to achieve 2°C suggest temperatures are unlikely to increase significantly faster than the 0.2°C per decade already observed between the 1990s and 2000s (Frame 2012; Rogelj *et al.* 2011b; Karl *et al.* 2000; Trenberth & Jones 2007). Hence, under the 2°C scenario, the peak rate of warming may well have already passed. Alternatively, under substantially higher emission scenarios, the peak warming rate is instead expected around the time of peak LLCF emissions (Smith *et al.* 2012), often in several decades time (Figure 7.7). Hence, the way to reduce the peak rate of warming using SLCF emissions is to reduce them gradually over the period in which LLCF emissions peak (e.g. dotted purple curve in Figure 7.7c), since this will reduce the peak rate of temperature change.

7.6.2 Implications

Several factors make immediate reductions in SLCFs attractive. There are substantial short-term co-benefits to agriculture, ecosystem services and human health (UNEP & WMO 2011; Shindell *et al.* 2012; Huntingford *et al.* 2011). Taking measures to reduce SLCF emissions soon will delay warming, and will reduce near-term rates of warming (Ramanathan & Xu 2010; UNEP & WMO 2011; Shindell *et al.* 2012). Early action may provide time to develop and implement SLCF reduction technologies such that they are available when needed, and some SLCF measures are argued to be cost-negative (UNEP & WMO 2011). Advocates of near-term SLCF reductions also argue that SLCF reductions may help build political momentum for reductions in CO_2, since the long timescales associated with the accumulation of LLCF emissions and with the emergence of their climate impacts present particular challenges to negotiators (Wallack & Ramanathan 2009; Victor *et al.* 2012; Barrett 2008; Bosetti & Victor 2011; Montzka *et al.* 2011; Rogelj *et al.* 2013). But it is important to be realistic about what near-term SLCF measures can actually achieve. It is tempting to claim that immediate action on SLCFs will contribute directly to a reduction in the overall risk of crossing potentially dangerous temperature thresholds. But unless we see a substantial turnaround in global CO_2 emission rates, such a claim would be untrue: if CO_2 emissions continue, near-term SLCF measures can delay the crossing of temperature thresholds, but they do not significantly mitigate the overall risk. Nor do they reduce the peak rate of warming, unless these SLCF measures are carefully timed to coincide with peak LLCF emissions.

In summary, by aiming for 2°C maximum global warming, governments implicitly committed to limiting total cumulative future emissions of LLCFs such as CO_2 (on one estimate to less than the equivalent of half a trillion tonnes of carbon (Allen *et al.* 2009c)). If emissions of LLCFs are about to peak and decline to meet this goal, then SLCFs will soon become a climate policy priority as a key determinant of peak warming. However if LLCF emissions continue to grow, SLCFs will remain largely irrelevant to the risk of cross-

ing potentially dangerous warming thresholds. Thus the effectiveness of SLCF reductions depends on the progress of LLCF reductions, while the reverse is not true. We argue LLCF reductions need to remain the *sine qua non* of mitigation policy.

7.7 Ongoing and future work

There are a number of ways that this work could be extended and improved, and some of them are explored below.

7.7.1 Improvements and verification

As in most of the chapters in this thesis, there are a number of improvements that could be made to the box-diffusion model, in particular to the representation of SLCFs . There is also work to be done in verifying the results presented here, both in other simple climate models such as MAGICC and also in more complex models such as EMICs and ESMs.

7.7.2 Optimal timing of SLCF cuts

In Section 7.6 I suggested that the optimal timing of SLCF cuts for reducing peak warming is after LLCF emissions peak but shortly before LLCF emissions go to zero. I also suggested that the optimal timing for SLCF cuts to reduce the peak warming rate might be just before the peak in LLCF emissions. It would be interesting to investigate these claims in more detail under a wider range of emissions pathways and show them with more certainty than simply by showing a few illustrative examples.

7.7.3 Metrics and SLCFs

At some point governments may have to choose how to allocate budgets between LLCFs and SLCFs , and there is still an open question as to how to do this. Much of this decision depends on how much we value different things such as air quality, the rate of warming, the magnitude of warming, etc. If all we consider for the moment is the magnitude of warming then we could establish a metric for trading between these gases based on their impact

on peak warming, otherwise known as GTP (Shine *et al.* 2005). We could calculate the GTPs of black carbon and tropospheric ozone, and we could also do it in the full likelihood framework as neither of these things have been done to date to the best of my knowledge (Shine *et al.* 2007).

7.7.4 Tipping points and SLCFs

Lenton *et al.* (2008) and others (Collins *et al.* 2013) have suggested that there may be non-linear *tipping points* in the climate system. Depending on how the triggers of these tipping points are defined we would want to cut SLCFs at different times to reduce the chances of reaching these tipping points. For example if the trigger is a certain temperature threshold, we might want to use SLCFs before or in the run up to the time of peak warming to avoid crossing the critical threshold. If the trigger is a rate of warming threshold then we might want to reduce SLCF emissions in the run up to and during the time of peak LLCF emissions. If the threshold was a 'degree-years' threshold, depending on the time spent above a temperature threshold, then we would want to use SLCFs before crossing the threshold and use carbon dioxide removal (CDR) or solar radiation management (SRM) (Vaughan *et al.* 2009) to come back down under the threshold as soon as possible after it was breached. These hypotheses could all be tested.

7.8 Conclusions

Throughout this chapter I have made the case that delaying reductions in SLCFs until CO_2 emissions start declining will not significantly impact on peak warming. This is not to say that we should wait to start emissions reductions, only that this point is often overlooked, for example by Molina *et al.* (2009) and Montzka *et al.* (2011).

The question of whether we should cut SLCFs now or later is more nuanced, and turns on factors such as cost, the short-term benefits for example to agriculture and health, and

political questions such as whether work to reduce SLCFs would make reductions in CO_2 more or less likely.

The basic physics of the central argument in this chapter is sound, namely that SLCFs have short lifetimes and so cannot have significant impacts on the climate in 100 years, other than through feedback effects. However the question of whether we can delay SLCF emissions reductions does depend on the shape of the baseline emissions profile. For example, for a profile that brought CO_2 emissions immediately to zero, delaying SLCF reductions would have a substantial impact on peak warming. If we only consider plausible emissions pathways that peak at or above 2°C, where we define plausible as CO_2 emissions not falling by more than 3% per year, then we are able to delay SLCF emissions reductions without substantially impacting peak warming. However, as discussed in Chapters 2 and 3, the box-diffusion model used to create these plots is very simple, and so this work would benefit from being reproduced in more complex climate models such as EMICs and ESMs.

This chapter was the final one in my thesis that contains new modelling results, so I will use the remaining chapter to conclude and discuss the implications of this work.

Chapter 8

Conclusions

8.1 Thesis aims

In this thesis I aimed to investigate cumulative emissions in order to refine their definition, and in order to help make the concept of cumulative carbon emissions more compatible with non-CO_2 GHGs.

I have proposed a definition of cumulative emissions of $\int_{t=\mathrm{yr}1750}^{t=\mathrm{yr}2200} E dt$ using emissions floors to refine the limits of the integral as discussed in Chapter 5. I have shown that cumulative carbon emissions under this definition provide a stronger constraint on CO_2-induced peak warming than both cumulative emissions under a variety of other definitions, and a number of alternative carbon emissions targets. I also showed that this property holds in the box-diffusion model at lower likelihoods.

Work to make the concept of cumulative emissions more compatible with non-CO_2 emissions using a combined metric was carried out in collaboration with Smith *et al.* (2012). I also suggested that delaying proposed reductions to SLCF emissions would have little impact on peak warming, and that these measures cannot be used to "buy time" to cut CO_2 emissions (UNEP & WMO 2011).

The findings in this thesis have been verified only in the box-diffusion model, which is outlined and tested here, and work remains to be done in exploring these results in other models.

In this chapter I present a number of recent results that have been produced using the

concept of cumulative emissions, before outlining a global climate policy that illustrates the uses of the various findings presented throughout this thesis.

8.2 Estimating the transient climate response to emissions

Over the course of writing this thesis, the concept of cumulative carbon emissions has received attention in a number of papers, and ultimately in the Working Group 1 (WG1) Summary for Policymakers (SPM) of the most recent IPCC report, which was published in the weeks leading up to the submission of this thesis (IPCC 2013). A key metric used in this report was transient climate response to emission (TCRE). TCRE is the TCR divided by the cumulative emissions from pre-industrial up until the time that the TCR is measured (Planton 2013; Gillett *et al.* 2013). Estimating TCRE was not an aim of this thesis, however the box-diffusion model does produce an estimate of TCRE, which can be compared with these other estimates in the IPCC report. This comparison has been done in Figure 8.1, which is taken from Chapter 12 of the latest IPCC report (Collins *et al.* 2013).

Because many of the plots in Figure 8.1 use different units and metrics for both temperature and emissions, we cannot do a direct comparison between some of these plots, but we can still estimate TCRE in many of these cases. The estimates discussed here are taken from Collins *et al.* (2013). In panel a Matthews *et al.* (2009) estimate the 5–95% range in TCRE as 1–2.1°C TtC^{-1}. In panel d, Johns *et al.* (2011) estimate a range in TCRE of 1–4°C TtC^{-1}. In panel e, Rogelj *et al.* (2012a) estimate a 5–95% range in TCRE of 1–2°C TtC^{-1}. Panel g is the box-diffusion model, and gives in a 5–95% confidence interval on the TCRE of 1.4–2.5°C TtC^{-1} from Allen *et al.* (2009c); the range in peak warming shown in the figure is 1.3–3.9°C TtC^{-1}. The difference between these two estimates from Allen *et al.* (2009c) is due to the warming that would occur after a complete cessation of emissions, which Allen *et al.* (2009c) refers to as the zero emissions commitment (ZEC). In panel h, Zickfeld *et al.* (2013) use EMICs to estimate a range in TCRE of 1.4–2.5°C TtC^{-1}, and a mean TCRE of 1.9°C TtC^{-1}. Collins *et al.* (2013) combines these estimates using expert

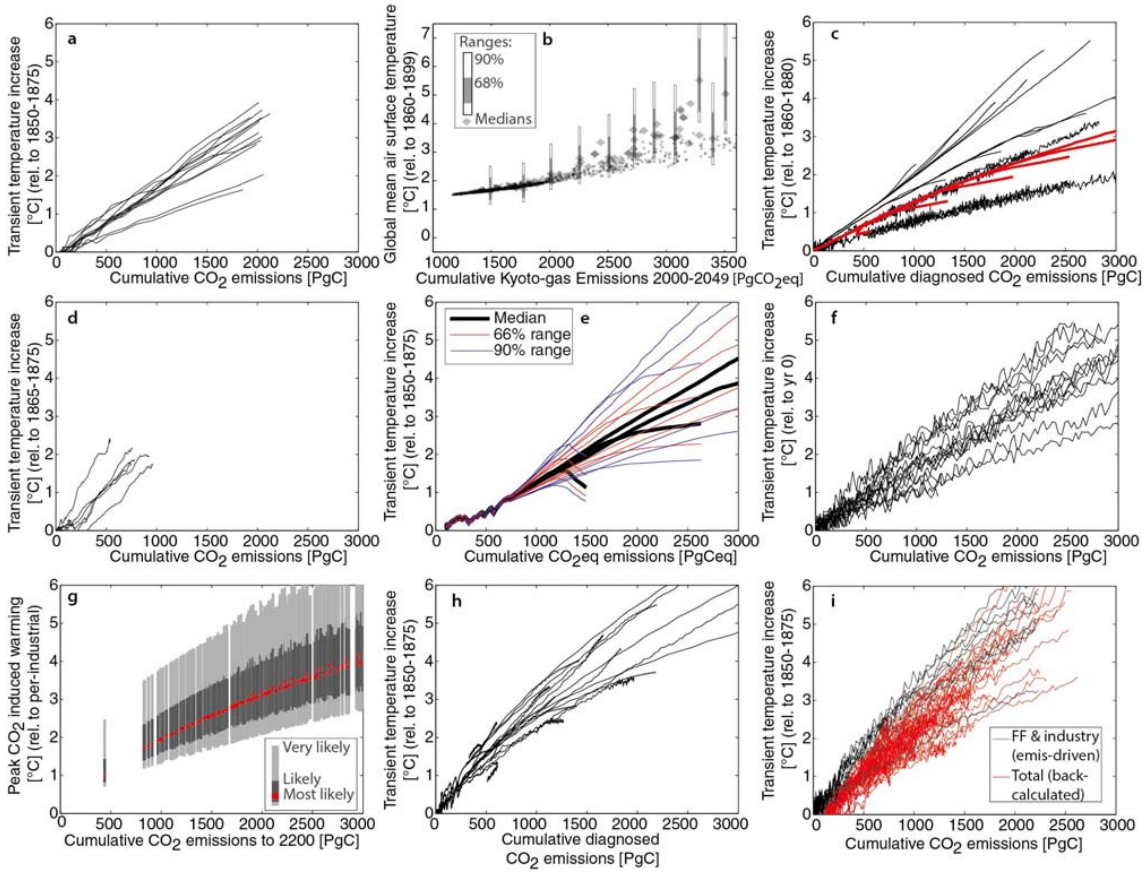


Figure 8.1: Global temperature change vs. cumulative carbon emissions for different scenarios and models. a) Transient global temperature increase vs. cumulative CO_2 emissions for C^4MIP Matthews *et al.* (2009), b) maximum temperature increase until 2100 vs. cumulative Kyoto-gas emissions (NB: this panel uses CO_2 equivalence, while all other panels are given in C equivalent.) (Meinshausen *et al.* 2009), c) transient temperature increase vs. cumulative CO2 emissions for IPCC TAR models (red) and IPCC AR4 EMICs (black), d) as in panel (a) but for the ENSEMBLES E1 scenario (Johns *et al.* 2011) e) transient temperature increase for the RCP scenarios based on the MAGICC (Rogelj *et al.* 2012a), f) transient temperature change from the CMIP5 1% yr^1 concentration driven simulations, g) peak CO_2-induced warming vs. cumulative CO_2 emissions to 2200 (adapted from my Figure 4.2), h) transient temperature increase from EMIC RCP simulations (Zickfeld *et al.* 2013) and i) transient temperature change from the CMIP5 historical and RCP8.5 emission driven simulations (black) and transient temperature change in all concentration-driven CMIP5 RCP simulations with back-calculated emissions (red). Source: IPCC AR5 WG1 Chapter 12 (Collins *et al.* 2013).

judgement to give a likely range of 0.8–2.5°C TtC^{-1} for cumulative emissions of less than approximately 2TtC until the time at which temperatures peak. Because Meinshausen *et al.* (2009) uses cumulative emissions until 2050 in panel b, we cannot directly calculate TCRE from this plot, as discussed in Chapter 4.

We find that the box-diffusion model's TCRE is within, but in the upper two-thirds of, the IPCC range. This may be in part due to the carbon cycle feedbacks within the model, as discussed in Chapter 3.

8.3 Designing a policy based on cumulative emissions

In this section I will assess the benefits and limitations of the work presented in this thesis by proposing a climate policy framework based on this work and discussing the benefits and limitations of this proposal.

8.3.1 Choosing a cumulative budget

This climate policy proposal is based around the concept of setting a cumulative emissions target, and then trading between gases using GTP as introduced in Section 3.5.2. Nations would begin by setting a temperature target, for example by agreeing to limit global mean temperature increase to 2°C above pre-industrial levels. This temperature target would then be used to calculate a cumulative emissions target using the TCRE estimates described in Section 8.2. A single cumulative carbon emissions target would be chosen, for example 0.8 TtC.

The choice of cumulative budget could change as we learn more about climate impacts, TCRE, and the cost of mitigation. This would be better from an economic standpoint, but also brings with it considerable credibility issues, namely a concern that policymakers would just increase the cumulative budget when there was political opposition to the target (Hepburn 2005). Some of this worry could be abated by setting up an independent body to update the target as we learn, similar to the UK Climate Change Committee (UK Commit-

tee on Climate Change 2008), however the international nature of this organisation could mean that it would not have enough credibility with national politicians (Hepburn 2005).

8.3.2 Choosing a framework

Having set a cumulative emissions target, nations would then decide upon a framework with which to apportion this cumulative budget between different nations, and ultimately different actors. I will discuss three possible proposals here which were mentioned briefly in Chapter 1, but in reality any number of proposals based on cumulative budgets could be used for this purpose.

One option is to do the apportioning 'upstream', i.e. at the point of extraction, using the Kyoto2 framework to auction permits to extract carbon out of the ground (Tickell 2008). Permit allocation could also be done using 'grandfathering' (i.e. based on current use), as a fraction of proven reserves, or by negotiation (Cramton & Kerr 2002).

Another upstream option would be to do this apportioning by regulating the fraction of carbon that has to be sequestered, α, to increase from zero today to approaching 100% as society approaches cumulative emissions target (Allen *et al.* 2009a). Under this 'SAFE Carbon' proposal, actors would trade SAFE carbon permits indicating the amount of carbon they had sequestered. They would then have to show regulators that after this trading, the ratio between their SAFE carbon permits and their carbon extracted was at least α.

A third option for apportioning this cumulative budget would be to use the Contraction and Convergence framework (Stott 2006). This cap-and-trade approach constrains the amount of carbon that each nation can emit, with the option of converging on equal per capita annual emissions allocations, before emissions allocations ultimately trend to near-zero (Stott 2006).

Several other frameworks have been proposed for apportioning these carbon emissions, from 'greenhouse development rights' to a global auction followed by trading under the cumulative cap (Baer *et al.* 2009; England *et al.* 2009). A discussion of the benefits and limitations of each of these frameworks is beyond the scope of this thesis, but choice of

framework is likely to be a trade-off between what is politically feasible and what is credible enough to ensure the cumulative budget is not exceeded, one would hope with the principles of equity and efficiency also in mind (Cramton & Kerr 2002; Baer *et al.* 2009).

8.3.3 Choosing a metric

Everything discussed thus far only applies to carbon emissions. Once the cumulative budget has been set and a policy framework has been chosen, I can now set up trading between CO_2 and non-CO_2 GHGs. As we are ultimately attempting to meet a temperature target I would chose to trade using a cost-effectiveness metric such as global cost potential (GCP) or GTP, as opposed to a cost-benefit metric such as global damages potential (GDP) or GWP (Tol *et al.* 2012). If we were to trade using either GCP or GTP, we would be trading gases according to their impact on peak warming (Tol *et al.* 2012; Shine *et al.* 2007). However GTP does not factor in the cost of abatement, which means that it is economically a sub-optimal policy compared with trading according to GCP (Tol *et al.* 2012). GCP sets up the trading prices of different gases across time so as to minimise the overall cost of abatement while still staying below a given temperature limit, resulting in a less expensive policy overall (Tol *et al.* 2012; Manne & Richels 2001). While GCP is the economically optimal metric to use, I am going to use GTP here because it is simpler, and because it does not rely on complex economic models in order to determine the prices of different gases at different times (Manne & Richels 2001). I am also choosing to use GTP in this chapter so that I can focus my discussion more on the physical aspects of this proposed policy, and less on its economic benefits and limitations.

The choice of time horizon in any cost-effectiveness metric is a crucial one, and I would choose for the time horizon to be to the time of peak warming (Shine *et al.* 2007). As an initial estimate I would consider using year 2200 as the time horizon, as my work in Chapter 5 suggests this is a physically-based time horizon over which to set a cumulative budget. Before embarking on this policy I would want to run a number of plausible emissions scenarios that keep temperatures below the temperature limit to find when temperature peak

in these model runs, and use this as the time horizon. I would update this time horizon as society gained more information about the response timescale of the climate system, TCRE, and in turn when peak warming would be likely to occur. I will call this metric global peak temperature change potential (GPTP), which is a specific case of GTP (Shine *et al.* 2007). I will return to this issue of time horizons in Section 8.3.5 on limitations, but first I will discuss some of the benefits of this proposal.

8.3.4 Benefits

Many of the benefits of this policy proposal arise from issues already discussed in this thesis.

8.3.4.1 Other types of emission targets

A cumulative carbon emissions target is a tighter constraint on peak warming than many other types of carbon emissions target. In Chapter 4 I showed that it is a tighter constraint than an carbon emissions target in year 2020 or 2050, or a target specifying the year of peak CO_2 emissions.

I also showed that the choice of time horizon, a cumulative target to the time of peak warming which occurs around 2200 in the box-diffusion model, is a better constraint on peak warming than cumulative emissions in 2050 as suggested by Meinshausen *et al.* (2009), or cumulative emissions to 2500 as suggested by Allen *et al.* (2009c). Unfortunately the time of peak warming is both model and scenario dependent, meaning that the exact timeframe over which the cumulative budget is set is poorly defined, as shown in Figure 6.7. I discuss this limitation more in Section 8.3.5.

8.3.4.2 Emission floors

By setting a cumulative emissions budget to the time of peak warming, we account for the fact that society may encounter one or more sectors that it is very difficult to decarbonise, a concept known as an emissions floor (see Chapter 5). In practise if society plans for an

emissions pathway without an emissions floor and then goes on to encounter one, society may find itself in a position where it is forced to exceed the cumulative budget. Though it is a very difficult question, more research into the question of whether we are likely to encounter an emissions floor could help avoid this problem. Alternatively society could simply adopt the precautionary principle and assume we will encounter an emissions floor following Anderson (2008), despite this approach's potential economic sub-optimality.

8.3.4.3 Approximate linearity in TCRE

As discussed in Chapter 4 and as shown in Figure 8.1, different models show differing degrees of linearity in their temperature response to cumulative emissions. The approach outlined here does not require for TCRE to be constant. By simply setting a target, with the option to adjust it in the future as we learn more about the implications of the target, we make no assumptions that TCRE is linear. We only use the fact that cumulative emissions is a strong constraint on future peak warming.

8.3.4.4 Non-CO_2 GHGs

One of the major benefits of this proposal over any other policy proposal that I have seen based on cumulative emissions is that it treats CO_2 and non-CO_2 GHGs in a unified framework, allowing them to be traded (Smith *et al.* 2012; Matthews *et al.* 2012).

In Smith *et al.* (2012), we created a framework that used cumulative emissions of long-lived climate forcers (LLCFs) and rates of emission of short-lived climate forcers (SLCFs). One of the problems with trading under this approach is that it would require trading of cumulative emissions (or stocks in economic parlance) of LLCFs with sustained rates of emissions (sustained flows in economic parlance). While this is possible in theory, it adds a complexity that has never been tested before in emissions trading markets. Using GTP requires trading gases at ratios that vary as a function of time, which to the best of my knowledge has not been tested before in international emissions trading markets, let alone adding

something as relatively complex as attempting to trade stocks against flows (Boucher & Reddy 2008).

Additionally, trading SLCFs as a flow, but trading based on each flow's impact on peak warming following Smith *et al.* (2012), means the relative ratios between the gases would be fixed. This would result in SLCFs being overpriced near the present day relative to what is optimal, because their prices should increase as we approach the time of peak warming if the aim is to avoid peak warming (Shine *et al.* 2007; Manne & Richels 2001), but instead they would be fixed.

Using GPTP also allows us to include carbon cycle feedbacks in our estimates of GPTP, which may make up a considerable component of the current GPTP of very short-lived species such as black carbon as discussed in Chapter 7. Current GWP-based emissions trading schemes do not include this carbon cycle feedback (Gillett & Matthews 2010).

8.3.5 Limitations

While this proposed climate policy based on cumulative emissions has some benefits, it also has a number of limitations.

8.3.5.1 Limitations of cumulative emissions

TCRE has a wide uncertainty range, as shown in Figures 1.3b, 4.1a and 8.1. The lower bound of the IPCC likely range in TCRE is more than three times larger than the upper bound of the IPCC likely range (Collins *et al.* 2013). This means that there is going to be a large inherent uncertainty in the warming due to any cumulative emissions target. Either this means that there is uncertainty over whether a cumulative budget will keep temperatures below a given temperature target, or it means that society must alter its cumulative budget as it learns more about TCRE (Allen & Frame 2007). Neither of these options are ideal, though this problem also affects any sort of emissions target due to to similar uncertainty in TCR and ECS (Huntingford *et al.* 2012; Collins *et al.* 2013).

Another problem with setting a cumulative emissions target is that there may be non-linearities in the climate system (Lenton *et al.* 2008; Collins *et al.* 2013). As discussed in Section 3.7.1 there are a number of components in the climate system that may cause the climate to exhibit highly non-linear behaviour. Thus far climate models have suggested that the climate behaves largely linearly below approximately 4°C of warming, (Matthews *et al.* 2009; Gillett *et al.* 2013), however this does not mean that this is how the climate will behave as these non-linearities may be poorly represented, if represented at all, in climate models (Collins *et al.* 2013). This means that temperatures could warm considerably more than anticipated for a given cumulative emissions budget.

The box-diffusion model, because it is a simple model, also does not include a number of key components of the climate system as discussed in Section 3.7.1. Many of these components are included in the models that the box-diffusion model is tuned to, but some are not. For example a reasonably accurate representation of ocean dynamics is not included in some of the C^4MIP models, meaning that it is possible that the box-diffusion model lacks some of the long-term memory that the oceans provide in the real world. This could result in a systematic error in given a cumulative emissions target.

Additionally the box-diffusion model does not include any dynamics, and so does not model the impact of inter-annual variability on global mean temperatures and the risk of exceeding temperature thresholds. This issue is discussed further in Section 3.7.1.

There are also a number of policy limitations to cumulative emissions which I discussed in Section 1.3.5.

8.3.5.2 Time horizons

Global peak temperature potential (GPTP), i.e. GTP to the time of peak warming, is both model- and scenario-dependent because the time of peak warming is both model- and scenario-dependent (Shine *et al.* 2007).

In this thesis I suggest that a cumulative carbon budget to 2200 would constrain CO_2-induced warming, but this does not mean that 2200 is the correct time horizon to use for

the GTP-based trading scheme. The time of peak warming is not necessarily near 2200, and in fact in many scenarios and models it is not, as shown in Figure 6.7. For example in MAGICC 6 under RCP-3PD temperatures peak in the mid- to late-21st Century.

The fact that models disagree considerably on the time of peak warming, suggesting that more modelling work would need to be done in order to establish when the time horizon would need to be in this policy. The time horizon would also depend on the temperature target set, and the approximate emissions pathways that we assume society would follow under that target (Shine *et al.* 2007).

This time-horizon dependence of this policy is potentially problematic because if the time horizon is selected incorrectly then the price of SLCFs will increase either too soon or too late, in the first case burdening society with unnecessary costs, and in the second case risking exceeding the temperature target (Manne & Richels 2001; Shine *et al.* 2007). This means that, like the cumulative budget, this time horizon may have to be shifted as we learn more about the climate system.

The benefit of shifting this time horizon is that it also helps balance emissions between short- and long-lived species through a positive feedback loop (where I use the term positive feedback loop here in the engineering and control theory sense where the feedback returns the system to equilibrium, not in the climate sense where a positive feedback drives the system away from equilibrium). If society starts emitting more SLCFs than expected, this will bring the estimate of the date of peak warming closer to present day, which increases the GPTP of these short-lived species and hence increases their price ratio and consequently decreases their emissions. This feedback mechanism also works the other way around to reduce the price ratio of SLCFs if less are being emitted than expected. This feedback mechanism helps ensure that a given temperature target is met, assuming all actors abide by the policy and the estimate of TCRE is accurate or is updated as we learn more about the climate system.

8.3.5.3 Other climate impacts

This policy proposal is predicated on the idea of avoiding a certain level of peak warming, or avoiding crossing a warming threshold. This would be an ideal basis for climate policy if all climate impacts scaled with the level of warming. The pattern scaling literature suggests that many climate impacts do scale linearly with the level of warming, however there are still several climate impacts that do not appear to, such as Arctic sea ice loss (Mitchell 2003; Comiso *et al.* 2008).

For example, Root *et al.* (2003), Leemans & Eickhout (2004), Jump & Penuelas (2005) and others have suggested that the rate of warming may control some potential climate impacts, such as the rate at which ecosystems need to adapt to warming. In Chapter 4 I suggested that the peak rate of CO_2-induced warming may be controlled by limiting the maximum rate of emission of CO_2. I went on to speculate that the derivative of the conclusion we reached in Smith *et al.* (2012) may also hold: namely that the rate of warming may be controlled by the rate of emission of LLCFs, and the rate of change in emissions of SLCFs.

Wu *et al.* (2010) suggests that some climate impacts, such as changes in global mean precipitation, may exhibit hysteresis. Van den Broeke *et al.* (2010) suggest that ice sheet melt may depend on the integral of temperature above some threshold with respect to time. Zickfeld *et al.* (2012) suggest that sea level rise shows scenario-dependence for a given level of warming despite the fact the cumulative emissions do not. These various impacts that are not linked directly to peak warming suggest that while cumulative emissions may constrain some climate impacts well, it cannot constrain all climate impacts well. Thus some climate impacts will either have to be less well constrained, or the climate policy will need to have additional targets added to it to constrain these quantitates (Steinacher *et al.* 2013).

8.3.5.4 Regional impacts

Though climate change is a global phenomenon, its impacts are felt locally. These impacts do not always scale directly with global mean temperatures as I discuss in Section 3.7.1. Some climate forcers such as CO_2 are well mixed, while others such as black carbon are not (Ramanathan & Carmichael 2008; Bond *et al.* 2013). Black carbon is emitted locally, and a significant portion of its climate impact can be felt locally (Shindell *et al.* 2012), as I discuss in Section 6.3.1. This suggests that regional mitigation plans could be used in addition to the policy suggested here to help avoid some of the worst impacts of regional climate change (Shindell *et al.* 2012).

8.3.5.5 Geoengineering

The technology behind geoengineering is not yet ready for adoption, and there is still considerably debate over whether it should even be researched (Bunzl 2009; Rayner *et al.* 2013). Nonetheless, this policy proposal ignores some of the potential of geoengineering.

The most obvious example of how geoengineering is not considered in this policy is solar radiation management (SRM). SRM is artificially altering the Earth's albedo, so as to reflect away more incoming solar radiation thereby reducing global warming (Vaughan & Lenton 2011). Proposed methods of doing this include injecting sulphate particles into the stratosphere, putting mirrors into orbit to block some incoming solar radiation, and directly reducing the albedo of the surface by covering darker areas with lighter materials, for example using white paint (Vaughan & Lenton 2011).

If SRM were adopted and considered safe then peak warming could be reduced below the levels otherwise expected without SRM (Wigley 2006; Vaughan & Lenton 2012). However SRM carries a termination risk if it is used to artificially suppress global temperature increases, as well as potential side-effects depending on how it is implemented, for example it would not help reduce ocean acidification (Ross & Matthews 2009; Vaughan & Lenton 2011). SRM could be incorporated into this proposal by setting the maximum

level of SRM in Wm^{-2} that is politically, technologically and economically feasible, and then increasing the cumulative budget so that the additional CO_2 concentrations and SRM match in terms of their equal by opposite contributions to radiative forcing (Wigley 2006; Vaughan & Lenton 2012).

Carbon dioxide removal (CDR) from the atmosphere, ocean or biosphere is another form of geoengineering that has not been fully accounted for in this policy proposal (Vaughan & Lenton 2011). When accounting for negative emissions, the emissions considered thus far need to be redefined as the *net* emission. Thus far in this policy proposal I have not accounted for the possibility of net negative carbon emissions. If the net negative emissions are after the time of peak warming then we have an 'overshoot' scenario in which temperatures reduce rapidly after they peak (Friedlingstein *et al.* 2011; van Vuuren *et al.* 2011; Parry *et al.* 2009). For simplicity's sake I will consider only what happens up until the time of peak warming in this policy proposal, and suggest that negative emissions could still be used in this proposal to reduce the size of a floor in emissions of any climate forcer.

8.3.6 Summary

The policy proposal I outline here is not a policy recommendation as it has considerable downsides associated with it, but it is a relatively novel proposal designed to highlight some of the benefits and limitations of my research in this thesis, while incorporating some recent developments in the field.

This proposal is built on my work in Chapter 4 showing that cumulative emissions targets provide a stronger constraint on CO_2-induced warming than other types of emissions targets. It also allows, at least in principle, for emissions floors, though in practise a large unexpected emissions floor would likely cause society to exceed its cumulative budget as discussed in Chapter 5. The policy allows for trading between CO_2 and non-CO_2 GHGs using GPTP as introduced in Chapters 3 and 8, and allows for a lower price ratio for SLCFs in the short-term than near the time of peak warming following the work of Manne & Richels (2001), Shine *et al.* (2007) and my work in Chapter 7.

This work was developed using the box-diffusion model introduced in Chapter 2, which has a different CO_2 concentration response function to a pulse of carbon emissions compared to most models, and possibly an overactive carbon cycle feedback, however it is still able to mimic the multi-model mean of the CMIP5 ensemble over the 21st Century, as discussed in Chapter 3. I also added seven non-CO_2 GHGs to the model using a single constant lifetime for each GHG, and a four-term radiative forcing equation. I used a shorter methane lifetime than recommended by the IPCC (Myhre & Shindell 2013) as it was a closer fit to historical observations as shown in Chapter 6 and matched to the results of Holmes *et al.* (2013). The absolute metrics for these gases behaved like those estimated by the IPCC, but the ratios relative to CO_2 were different because of the box-diffusion models' different impulse response (Myhre & Shindell 2013). If the policy proposal suggested here were to be taken forward it would need to be tested using additional models, particularly the component relying on predicting the time at which temperatures peak which is both model- and scenario-dependent.

8.4 Comparison with other work

Having discussed the work in this thesis through the lens of a policy proposal, I will now briefly discuss it relative to similar work.

8.4.1 Emissions targets

In this thesis I ask which type of emissions target best constrains peak warming. Work has been done in the past on emissions targets and their relationship to peak warming. Rogelj *et al.* (2011b) and Rogelj *et al.* (2012b) use emissions corridors to map the range of possible emissions rates in each year that give certain probabilities of exceeding 2°C. Rogelj *et al.* (2011a), Huntingford *et al.* (2012) and Rogelj *et al.* (2012b) look at the range of emissions in years 2020 and 2050 and consider their associated risk of exceeding temperature thresholds. In Huntingford *et al.* (2012) we also consider how the year at which emissions

peak and the subsequent rate of emissions decline and how this constrains the probability of exceeding 2°C. A number of reports from international agencies and NGOs also discuss this topic (UNEP 2012; Füssler *et al.* 2012).

A number of papers also discuss the relationship between cumulative emissions and peak warming (Allen *et al.* 2009c; Meinshausen *et al.* 2009; Matthews *et al.* 2009; Zickfeld *et al.* 2009; Raupach *et al.* 2011; Zickfeld *et al.* 2013; Gillett *et al.* 2013). In Chapter 4 I compare a number of these approaches to show that cumulative emissions provide a stronger constraint on CO_2-induced peak warming than emissions targets in a specific year, or peak year of emissions or the subsequent rate of emissions decline.

8.4.2 Emission floors

An initial worry with the concept of cumulative carbon emissions was that it might be incompatible with the concept of an emissions floor. Allen *et al.* (2009c) and Anderson (2008) only use emissions pathways that asymptote to zero, meaning that it is possible to integrate the entire emissions pathways to find a cumulative total. In Chapter 5 I considered whether cumulative emissions were compatible with an emissions floor, and found that they were. Even among emissions pathways with emissions floors, cumulative emissions did correlate well with peak warming.

A second question arising from the literature is what the appropriate integral on cumulative carbon emissions is. Meinshausen *et al.* (2009) considers cumulative emissions to 2050, Anderson (2008) and Anderson *et al.* (2008) consider cumulative emissions to 2100, Matthews *et al.* (2009) suggest it is cumulative emissions to the time of peak warming that controls peak warming, while Allen *et al.* (2009c) integrate cumulative emissions to 2500. In Chapters 4 and 5 I consider these various approaches and find that in the box-diffusion model it is cumulative emissions to approximately 2200 that provide the best constrain on peak warming, particularly when emissions floors are considered. I recognise that this result may be model dependent, and suggest that this conclusion is tested in other models before the result can be considered generalisable.

8.4.3 Short-lived climate forcers

In the past few years a number of papers have been published advocating near-term reductions in short-lived climate forcers (SLCFs). Most of these papers use co-benefits and reductions in near-term warming to motivate emissions reductions (UNEP & WMO 2011; Shindell *et al.* 2012; Ramanathan & Victor 2010; Victor *et al.* 2012; Ramanathan & Xu 2010). However two of these papers suggest that these reductions should be made in the near-term in order to reduce peak forcing (Montzka *et al.* 2011) and to reduce the probability of crossing temperature thresholds (Molina *et al.* 2009). These papers implicitly assume sustained emissions reductions and so find that near-term reductions would reduce forcings and warming levels in several decades, however they do not consider the impact of delaying these emissions reductions. I show in Chapter 7 that delaying emissions reductions in SLCFs does not significantly increase peak warming, because of the short lifetimes of SLCFs.

8.5 Policy implications

In addition to the specific policy proposal discussed in Section 8.3, some of the results in thesis thesis have further policy relevance. The results presented on emissions targets, cumulative emissions and emissions floors are relevant to the setting of both national and global emissions targets, suggesting that a cumulative carbon emissions target with appropriate limits on the integration could provide a stronger constraint on peak warming than existing policy proposals based on annual targets. I have presented these findings at the UK Department for Energy and Climate Change, UK HM Treasury, and in the UK House of Lords to the Labour Climate Change Spokesperson.

The work suggesting that SLCF reductions could be delayed without significant impact on peak warming has relevance for proposals being put forward internationally by the Climate and Clean Air Coalition of nations and championed by the US Department of State. This work does not suggest that SLCF reductions should not be made in the near-term,

only that reducing peak warming cannot be used as a motivation for these reductions being made now rather than later. These findings have been presented to UK HM Treasury and to the UNFCCC Subsidiary Body for Scientific and Technological Advice (SBSTA).

8.6 Conclusion

In this thesis I have attempted to extend and refine the concept of cumulative emissions to define its integral more precisely, to include emission floors in the concept, and to include non-CO_2 forcers. I found that in the box-diffusion model cumulative carbon emissions to 2200 provide a stronger constraint on peak CO_2-induced warming than any other type of emissions target currently proposed. I also showed that delaying near-term SLCF reductions does not have a significant impact on peak warming.

I have discussed possible avenues for future work in individual chapters, but it is clear that this work requires further testing in additional models, particularly the work on emissions floors which may well be model-dependent.

In writing this thesis, I hope I have contributed to our understanding of emission targets for avoiding dangerous climate change.

References

Allen, M. R. & D. J. Frame. Call Off the Quest. *Science*, **318**(5850):582–583 (2007).

Allen, M. R., D. J. Frame, & C. F. Mason. The case for mandatory sequestration. *Nature Geoscience*, **2**(12):813–814 (2009a).

Allen, M. R., *et al.* The exit strategy. *Nature Reports Climate Change* (2009b).

Allen, M. R., *et al.* Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, **458**(7242):1163–1166 (2009c).

Alley, R. B. Abrupt Climate Change. *Science*, **299**(5615):2005–2010 (2003).

Ammann, C. M. Statistical analysis of tropical explosive volcanism occurrences over the last 6 centuries. *Geophysical Research Letters*, **30**(5):1210 (2003).

Anderson, K. Reframing the climate change challenge in light of post-2000 emission trends. *Philosophical Transactions of the Royal Society A* (2008).

Anderson, K. & A. Bows. Beyond 'dangerous' climate change: emission scenarios for a new world. *Philosophical Transactions of the Royal Society A*, **369**(1934):20–44 (2011).

Anderson, K., A. Bows, & S. Mander. From long-term targets to cumulative emission pathways: Reframing UK climate policy. *Energy Policy*, **36**(10):3714–3722 (2008).

Andronova, N. G. & M. E. Schlesinger. Objective estimation of the probability density function for climate sensitivity. *Journal of Geophysical Research*, **106**(D19):22605 (2001).

4136 Archer, D. Fate of fossil fuel CO_2 in geologic time. *Journal of Geophysical Research*,
4137 **110**(C9):C09S05 (2005).

4138 Archer, D. Methane hydrate stability and anthropogenic climate change. *Biogeosciences*
4139 *Discussions*, **4**(2):993–1057 (2007).

4140 Archer, D. & V. Brovkin. The millennial atmospheric lifetime of anthropogenic CO_2.
4141 *Climatic Change*, **90**:283–297 (2008).

4142 Arneth, A., *et al.* Terrestrial biogeochemical feedbacks in the climate system. *Nature*
4143 *Geoscience*, **3**(8):525–532 (2010).

4144 Arora, V. K., *et al.* Carbon emission limits required to satisfy future representative concen-
4145 tration pathways of greenhouse gases. *Geophysical Research Letters*, **38** (2011).

4146 Arora, V. K., *et al.* Carbon-concentration and carbon-climate feedbacks in CMIP5 Earth
4147 system models. *Journal of Climate* (2013).

4148 Asrar, G. & J. Hurrell, editors. *Climate Science for serving society: research, modelling*
4149 *and prediction priorities*, chapter Atmospheric composition, irreversible climate change,
4150 and mitigation policy (Springer, 2012).

4151 Baede, A. P. M. *Climate Change 2007: The Physical Science Basis. Contribution of Work-*
4152 *ing Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate*
4153 *Change*, chapter Annex 1 (Cambridge University Press, 2007).

4154 Baer, P., *et al.* Greenhouse Development Rights: A Proposal for a Fair Global Climate
4155 Treaty. *Ethics, Place & Environment*, **12**(3):267–281 (2009).

4156 Baker, M. B. Cloud Microphysics and Climate. *Science*, **276**(5315):1072–1078 (1997).

4157 Barrett, S. Climate treaties and the imperative of enforcement. *Oxford Review of Economic*
4158 *Policy*, **24**(2):239–258 (2008).

Biello, D. How to buy time in the fight against climate change: Mobilize to stop soot and methane. *Scientific American* (2012).

Blundell, S. & K. Blundell. *Concepts in Thermal Physics* (Oxford University Press, 2006).

Boden, T., B. Andres, & G. Marland. Global co2 emissions from fossil-fuel burning, cement manufacture, and gas flaring: 1751-2010. Carbon Dioxide Information Analysis Center (CDAIC) (2013).

Bond, T. C., *et al.* Bounding the role of black carbon in the climate system: A scientific assessment. *Journal of Geophysical Research: Atmospheres*, **118**:1–173 (2013).

Bosetti, V. & D. G. Victor. Politics and economics of second-best regulation of greenhouse gases: The importance of regulatory credibility. *Energ J*, **32**(1):1–24 (2011).

Boucher, O. & D. Randall. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Clouds and Aerosols (Cambridge University Press, 2013).

Boucher, O. & M. S. Reddy. Climate trade-off between black carbon and carbon dioxide emissions. *Energy Policy*, **36**(1):193–200 (2008).

Boucher, O., *et al.* The indirect global warming potential and global temperature change potential due to methane oxidation. *Environmental Research Letters*, **4**(4):044007 (2009).

Boucher, O., *et al.* Reversibility in an Earth System model in response to CO_2 concentration changes. *Environmental Research Letters*, **7**(2):024013 (2012).

Bousquet, P., *et al.* Contribution of anthropogenic and natural sources to atmospheric methane variability. *Nature*, **443**(7110):439–443 (2006).

Bowerman, N. H. A., *et al.* Cumulative carbon emissions, emissions floors and short-term rates of warming: implications for policy. *Philos T R Soc A*, **369**(1934):45–66 (2011).

Bowerman, N. H. A., *et al.* *Can we still avoid dangerous climate change?*, chapter The importance of limiting cumulative emissions (Springer, 2013a).

Bowerman, N. H. A., *et al.* Short-lived climate pollutants and the two-degree goal. *Accepted* (2013b).

Box, G. E. P. Science and statistics. *Journal of the American Statistical Association* (1976).

Brier, G. W. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, **78**(1):1–3 (1950).

Brohan, P., *et al.* Uncertainty estimates in regional and global observed temperature changes: A new data set from 1850. *Journal of Geophysical Research*, **111** (2006).

Brovkin, V., *et al.* Sensitivity of a coupled climate-carbon cycle model to large volcanic eruptions during the last millennium. *Tellus B* (2010).

Budyko, M. I. The effect of solar radiation variations on the climate of the earth. *Tellus*, **21**(5):611–619 (1969).

Bunzl, M. Researching geoengineering: should not or could not? *Environmental Research Letters*, **4**(4):045104 (2009).

Butchart, N. & A. A. Scaife. Removal of chlorofluorocarbons by increased mass exchange between the stratosphere and troposphere in a changing climate. *Nature*, **410**(6830):799–802 (2001).

Canadell, J. G., *et al.* Contributions to accelerating atmospheric CO_2 growth from economic activity, carbon intensity, and efficiency of natural sinks. *Proceedings of the National Academy of Sciences* (2007).

Capotondi, A., *et al.* Enhanced upper ocean stratification with climate change in the CMIP3 models. *Journal of Geophysical Research: Atmospheres*, **117**(C4) (2012).

Christensen, J. H. & K. K. Kanikicharla. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Climate Phenomena and their Relevance for Future Regional Climate Change (Cambridge University Press, 2013).

Ciais, P. & C. Sabine. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Carbon and Other Biochemical Cycles (Cambridge University Press, 2013).

Claussen, M., *et al.* Did humankind prevent a Holocene glaciation? *Climatic Change*, **69**(2-3):409–417 (2005).

Collins, M., *et al. Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Long-term Climate Change: Projections, Commitments and Irreversibility (Cambridge University Press, 2013).

Collins, W. J., *et al.* Development and evaluation of an Earth-system model – HadGEM2. *Geoscientific Model Development Discussions*, **4**(2):997–1062 (2011).

Colman, R. A. & B. J. McAvaney. On tropospheric adjustment to forcing and climate feedbacks. *Climate Dynamics*, **36**:1649–1658 (2011).

Comiso, J. C. & D. G. Vaughan. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Observations: Cryosphere (Cambridge University Press, 2013).

Comiso, J. C., *et al.* Accelerated decline in the Arctic sea ice cover. *Geophysical Research Letters*, **35**(1):L01703 (2008).

Cox, P. & D. Stephenson. Climate change - a changing climate for prediction. *Science*, **317**(5835):207–208 (2007).

Cox, P. M. & H. A. Jeffery. Methane radiative forcing controls the allowable CO_2 emissions for climate stabilization. *Current Opinion in Environmental Sustainability*, **2**:404–408 (2010).

Cox, P. M., *et al.* Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature*, **408**(6809):184–187 (2000).

Cramton, P. & S. Kerr. Tradeable carbon permit auctions: How and why to auction not grandfather. *Energy Policy*, **30**(4):333–345 (2002).

Crank, J. The mathematics of diffusion. *Clarendon Press* (1979).

Crook, J. A. & P. M. Forster. A balance between radiative forcing and climate feedback in the modeled 20th century temperature response. *Journal of Geophysical Research*, **116**(D17):D17108 (2011).

Davis, S. J., K. Caldeira, & H. D. Matthews. Future CO_2 emissions and climate change from existing energy infrastructure. *Science*, **329**(5997):1330–1333 (2010).

Delworth, T. L. & T. R. Knutson. Simulation of early 20^{th} century global warming. *Science*, **287**(5461):2246–2250 (2000).

den Elzen, M. & M. Meinshausen. Meeting the EU 2 °C climate target: Global and regional emission implications. *Climate Policy*, **6**(5):545–564 (2006).

den Elzen, M., M. Meinshausen, & D. P. van Vuuren. Multi-gas emission envelopes to meet greenhouse gas concentration targets: Costs versus certainty of limiting temperature increase. *Global Environmental Change*, **17**(2):260–280 (2007).

Denman, K., *et al.* *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Couplings Between Changes in the Climate System and Biogeochemistry (Cambridge University Press, 2007).

Dolman, A. J., *et al.* A carbon cycle science update since IPCC AR4. *AMBIO*, **39**(5-6):402–412 (2010).

Domingues, C. M., *et al.* Improved estimates of upper-ocean warming and multi-decadal sea-level rise. *Nature*, **453**:1090–1093 (2008).

Dommenget, D. The Ocean's Role in Continental Climate Variability and Change. *Journal of Climate*, **22**(18):4939–4952 (2009).

Dyson, F. A meeting with Enrico Fermi. *Nature*, **427**:297 (2004).

Easterling, W. & P. Aggarwal. *Climate Change 2007: Impacts, Adaptation and Vulnerability*, chapter Food, fibre and forest products (Cambridge University Press, 2007).

Eby, M., *et al.* Lifetime of anthropogenic climate change: Millennial time scales of potential CO_2 and surface temperature perturbations. *Journal of Climate*, **22**:2501–2511 (2009).

Eby, M., *et al.* Historical and idealized climate model experiments: An EMIC intercomparison. *Submitted* (2013).

Eisenman, I. & J. S. Wettlaufer. Nonlinear threshold behavior during the loss of Arctic sea ice. *Proceedings of the National Academy of Sciences*, **106**(1):28–32 (2009).

England, M. H., A. S. Gupta, & A. J. Pitman. Constraining future greenhouse gas emissions by a cumulative target. *Proceedings of the National Academy of Sciences of the United States of America*, **106**(39):16539–16540 (2009).

Erikson, E. Air-ocean-icecap interactions in relations to climatic fluctuations and glaciation cycles. *Meteorological Monographs*, **8**:68–92 (1968).

Etheridge, D. M., L. P. Steele, & R. L. Langenfelds. Historical CO_2 records from the Law Dome DE08, DE08-2, and DSS ice cores. *Carbon Dioxide Information Analysis Center (CDIAC). Trends: A compendium of data on global change* (1998a).

Etheridge, D. M., *et al.* Natural and anthropogenic changes in atmospheric co_2 over the last 1000 years from air in Antarctic ice and firn. *Journal of Geophysical Research*, **101**(D2):4115 (1996).

Etheridge, D. M., *et al.* Atmospheric methane between 1000 A.D. and present: Evidence of anthropogenic emissions and climatic variability. *Journal of Geophysical Research*, **103**(D13):15979 (1998b).

Fankhauser, S., *et al. Climate Change 2001: Impact, Adaptation and Vulnerability. Contribution of Working Group II to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Vulnerability to Climate Change and Reasons for Concern: A Synthesis, pages 417–470 (Cambridge University Press, 2001).

Fischer, E. M., *et al.* European climate response to tropical volcanic eruptions over the last half millennium. *Geophysical Research Letters*, **34**(5):L05707 (2007).

Flato, G. & J. Marotzke. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Evaluation of Climate Models (Cambridge University Press, 2013).

Forster, P., K. P. Shine, & N. Stuber. It is premature to include non-CO_2 effects of aviation in emission trading schemes. *Atmospheric Environment*, **40**(6):1117–1121 (2006).

Forster, P., *et al.* *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Changes in Atmospheric Constituents and in Radiative Forcing (Cambridge University Press, 2007).

Frame, D., *et al.* Constraining climate forecasts: The role of prior assumptions. *Geophysical Research Letters*, **32**(9):L09702 (2005).

Frame, D. J. Assessment of the first consensus prediction on climate change. *Nature Climate Change*, **3**(4):357–359 (2012).

Frame, D. J., *et al.* The climateprediction.net BBC climate change experiment: design of the coupled model ensemble. *Philosophical Transactions of the Royal Society A*, **367**(1890):855–870 (2009).

Friedlingstein, P. Contributions of past and present human generations to committed warming caused by carbon dioxide. *Proceedings of the National Academy of Sciences*, **102**(31):10832–10836 (2005).

Friedlingstein, P., *et al.* Positive feedback between future climate change and the carbon cycle. *Geophysical Research Letters*, **28**(8):1543–1546 (2001).

Friedlingstein, P., *et al.* How positive is the feedback between climate change and the carbon cycle? *Tellus B*, **55**(2):692–700 (2003).

Friedlingstein, P., *et al.* Climate-carbon cycle feedback analysis: Results from the C^4MIP model intercomparison. *Journal of Climate*, **19**(14):3337–3353 (2006).

Friedlingstein, P., *et al.* Update on CO_2 emissions. *Nature Geoscience*, **3**(12):811–812 (2010).

Friedlingstein, P., *et al.* Long-term climate implications of twenty-first century options for carbon dioxide emission mitigation. *Nature Climate Change*, **1**(9):457–461 (2011).

Frölicher, T. L. & F. Joos. Reversible and irreversible impacts of greenhouse gas emissions in multi-century projections with the NCAR global coupled carbon cycle-climate model. *Climate Dynamics*, **35**(7-8):1439–1459 (2010).

Frölicher, T. L., F. Joos, & C. C. Raible. Sensitivity of atmospheric CO_2 and climate to explosive volcanic eruptions. *Biogeosciences* (2011).

Frölicher, T. L., *et al.* Atmospheric CO_2 response to volcanic eruptions: The role of ENSO, season, and variability. *Global Biogeochemical Cycles*, **27**:239–251 (2013).

Fuglestvedt, J. S., *et al.* Metrics of climate change: Assessing radiative forcing and emission indices. *Climatic Change*, **58**(3):267–331 (2003).

Fuglestvedt, J. S., *et al.* Transport impacts on atmosphere and climate: Metrics. *Atmospheric Environment*, **44**(37):4648–4677 (2010).

Fuss, S., *et al.* Impact of climate policy uncertainty on the adoption of electricity generating technologies. *Energy Policy*, **37**(2):733–743 (2009).

Füssler, J., *et al.* Emission pathways to reach 2°c target. Swiss Federal Office for the Environment (FOEN) (2012).

Fyke, J. G. & A. J. Weaver. The effect of potential future climate change on the marine methane hydrate stability zone. *Journal of Climate*, **19**(22):5903–5917 (2006).

G8. Declaration of the leaders: the Major Economies Forum on energy and climate (2009).

Gillett, N. P. & H. D. Matthews. Accounting for carbon cycle feedbacks in a comparison of the global warming effects of greenhouse gases. *Environmental Research Letters*, **5**(3) (2010).

Gillett, N. P., *et al.* Ongoing climate change following a complete cessation of carbon dioxide emissions. *Nature Geoscience*, **4**(2):83–87 (2011).

Gillett, N. P., *et al.* Constraining the ratio of global warming to cumulative CO_2 emissions using CMIP5 simulations. *Journal of Climate*, **26**:6844–2858 (2013).

Gleckler, P. J., K. E. Taylor, & C. Doutriaux. Performance metrics for climate models. *Journal of Geophysical Research*, **113**(D6):D06104 (2008).

Godal, O. & J. Fuglestvedt. Testing 100-year global warming potentials: impacts on compliance costs and abatement profile. *Climatic Change*, **52**(1-2):93–127 (2002).

Good, P., *et al.* Quantifying Environmental Drivers of Future Tropical Forest Extent. *Journal of Climate*, **24**(5):1337–1349 (2011).

Goody, R., J. Anderson, & G. North. Testing Climate Models: An Approach. *Bulletin of the American Meteorological Society*, **79**(11):2541–2549 (1998).

Govindasamy, B., *et al.* Limitations of the equivalent CO_2 approximation in climate change simulations. *Journal of Geophysical Research*, **106** (2001).

Gregory, J. M. Vertical heat transports in the ocean and their effect on time-dependent climate change. *Climate Dynamics*, **16**(7):501–515 (2000).

Guemas, V., *et al.* Retrospective prediction of the global warming slowdown in the past decade. *Nature Climate Change*, **3**(7):649–653 (2013).

Hammitt, J. K., *et al.* A welfare-based index for assessing environmental effects of greenhouse-gas emissions. *Nature*, **381**(6580):301–303 (1996).

Hansen, J., *et al.* Efficacy of climate forcings. *Journal of Geophysical Research*, **110** (2005).

Hansen, J., *et al.* Target atmospheric CO_2: Where should humanity aim? *Open Atmospheric Science Journal* (2008).

Hansen, J., *et al.* Global surface temperature change. *Reviews of Geophysics*, **48**(4):RG4004 (2010).

Hare, B. & M. Meinshausen. How much warming are we committed to and how much can be avoided? *Climatic Change*, **75**(1-2):111–149 (2006).

Hartmann, D. L., A. K. Tank, & M. Rusticucci. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Observations: Atmosphere and Surface (Cambridge University Press, 2013).

Harvey, L. D. D. A guide to global warming potentials (GWPs). *Energy Policy*, **21**(1):24–34 (1993).

Hawkins, E. & R. Sutton. The potential to narrow uncertainty in regional climate predictions. *Bulletin of the American Meteorological Society*, **90**(8) (2009).

Hepburn, C. *Climate-change policy*, chapter Credible carbon policy (Oxford University Press, 2005).

Hepburn, C. The energy mix, carbon pricing and border carbon adjustments. *Environmental Law & Management*, **24** (2012).

Höhne, N., *et al.* Contributions of individual countries' emissions to climate change and their uncertainty. *Climatic Change*, **106**(3):359–391 (2010).

Höhne, N., *et al.* National GHG emissions reduction pledges and 2 °C: comparison of studies. *Climate Policy*, **12**(3):356–377 (2012).

Holmes, C. D., *et al.* Future methane, hydroxyl, and their uncertainties: key climate and emission parameters for future predictions. *Atmospheric Chemistry and Physics*, **13**(1):285–302 (2013).

Houghton, R. A. Carbon flux to the atmosphere from land-use changes: 1850-2005. In *TRENDS: A Compendium of Data on Global Change* (Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tenn., U.S.A., 2008).

Houghton, R. A. How well do we know the flux of CO_2 from land-use change? *Tellus B* (2010).

Hua, Q. & M. Barbetti. Review of tropospheric bomb (super 14) C data for carbon cycle modeling and age calibration purposes. *Radiocarbon*, **46**(3):1273–1298 (2004).

Huang, B. Sensitivities of deep-ocean heat uptake and heat content to surface fluxes and subgrid-scale parameters in an ocean general circulation model with idealized geometry. *Journal of Geophysical Research*, **108**(C1):3015 (2003).

Huntingford, C., *et al.* Imogen: an intermediate complexity model to evaluate terrestrial impacts of a changing climate. *Geosci Model Dev*, **3**(2):679–687 (2010).

Huntingford, C., *et al.* Highly contrasting effects of different climate forcing agents on terrestrial ecosystem services. *Philos T R Soc A*, **369**(1943):2026–2037 (2011).

Huntingford, C., *et al.* The link between a global 2 degrees C warming threshold and emissions in years 2020, 2050 and beyond. *Environ Res Lett*, **7**(1):014039 (2012).

Huntingford, C., *et al.* Carbon dioxide and methane exchange implications under two metrics for heavy mitigation RCP2.6 scenario. *Submitted* (2013).

IPCC. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Summary for Policymakers (Cambridge University Press, 2007).

IPCC. *Special Report on Renewable Energy Sources and Climate Change Mitigation. Prepared by Working Group III of the Intergovernmental Panel on Climate Change* (Cambridge University Press, 2011).

IPCC. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Summary for Policymakers (Cambridge University Press, 2013).

Jewson, S., D. Rowlands, & M. R. Allen. A new method for making objective probabilistic climate forecasts from numerical climate models based on Jeffrey's prior. *arXiv:0908.4207* (2010a).

Jewson, S., D. Rowlands, & M. R. Allen. Objective climate model predictions using Jeffreys' prior: the general multivariate normal case. *arXiv:1005.3907* (2010b).

Jewson, S., D. Rowlands, & M. R. Allen. Objective probabilistic forecasts of future climate based on Jeffreys' prior: the case of correlated observables. *arXiv:1005.2354* (2010c).

Johansson, D. J. A. Economics- and physical-based metrics for comparing greenhouse gases. *Climatic Change*, **110**(1-2):123–141 (2011).

Johansson, D. J. A., U. M. Persson, & C. Azar. The cost of using global warming potentials: Analysing the trade off between CO_2, CH_4 and N_2O. *Climatic Change*, **77**(3-4):291–309 (2006).

Johns, T. C., *et al.* Climate change under aggressive mitigation: the ENSEMBLES multi-model experiment. *Climate Dynamics* (2011).

Jones, C., *et al.* Committed terrestrial ecosystem changes due to climate change. *Nature Geoscience*, **2**(7):484–487 (2009).

Jones, C. D. & P. M. Cox. Modeling the volcanic signal in the atmospheric CO_2 record. *Global Biogeochemical Cycles*, **15**(2):453–465 (2001).

Jones, C. D., P. M. Cox, & C. Huntingford. *Avoiding Dangerous Climate Change*, chapter Impact of climate–carbon cycle feedbacks on emissions scenarios to achieve stabilisation, pages 323–31 (Cambridge University Press, Cambridge, 2006).

Jones, P. D., *et al.* Hemispheric and largescale landsurface air temperature variations: An extensive revision and an update to 2010. *Journal of Geophysical Research: Atmospheres (1984–2012)*, **117**(D5) (2012).

Joos, F., *et al.* An efficient and accurate representation of complex oceanic and biospheric models of anthropogenic carbon uptake. *Tellus B*, **48**(3):397–417 (1996).

Joos, F., *et al.* Global warming feedbacks on terrestrial carbon uptake under the intergovernmental panel on climate change (IPCC) emission scenarios. *Global Biogeochemical Cycles*, **15**(4):891–907 (2001).

Joos, F., *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**(5):2793–2825 (2013).

Joshi, M., *et al.* Mechanisms for the land/sea warming contrast exhibited by simulations of climate change. *Climate Dynamics*, **30**(5):455–465 (2007).

Jump, A. & J. Penuelas. Running to stand still: adaptation and the response of plants to rapid climate change. *Ecol Lett*, **8**(9):1010–1020 (2005).

Kahneman, D. & A. Tversky. Intuitive Prediction: Biases and Corrective Procedures. *TIMS Studies in Management Science*, **12**:313–327 (1977).

Kallbekken, S., *et al.* Curbing emissions: cap and rate. *Nature Reports Climate Change* (2009).

Karl, T. R., R. W. Knight, & B. Baker. The record breaking global temperatures of 1997 and 1998: Evidence for an increase in the rate of global warming? **27**(5):719–722 (2000).

Kawai, Y. & A. Wada. Diurnal sea surface temperature variation and its impact on the atmosphere and ocean: A review. *Journal of Oceanography*, **63**(5):721–744 (2007).

Keeling, R. F., *et al.* Atmospheric carbon dioxide record from Mauna Loa. Carbon Dioxide Information Analysis Center (CDAIC) (2013).

Keith, D., M. Ha-Duong, & K. Stolaroff. Climate strategy with CO_2 capture from the air. *Climatic Change*, **74**:17–45 (2006).

Keller, K., D. McInerney, & D. Bradford. Carbon dioxide sequestration: how much and when? *Climatic Change*, **88**:267–291 (2008).

Kennedy, J. J., *et al.* Reassessing biases and other uncertainties in sea surface temperature observations measured in situ since 1850: 2. Biases and homogenization. *Journal of Geophysical Research: Atmospheres (1984–2012)*, **116**(D14) (2011).

Kim, J., *et al.* A seasonal precipitation and stream flow hindcast and prediction study in the western united states during the 1997/98 winter season using a dynamic downscaling system. *Journal of Hydrometeorology*, **1**(4):311–329 (2000).

Klein, R. J. & S. Huq. *Climate Change 2007: Impacts, Adaptation and Vulnerability*, chapter Inter-relationships between adaptation and mitigation (Cambridge University Press, 2007).

Kosaka, Y. & S.-P. Xie. Recent global-warming hiatus tied to equatorial Pacific surface cooling. *Nature*, **501**(7467):403–407 (2013).

Koven, C. D., *et al.* Permafrost carbon-climate feedbacks accelerate global warming. *Proceedings of the National Academy of Sciences* (2011).

Kug, J.-S., J.-Y. Lee, & I.-S. Kang. Global Sea Surface Temperature Prediction Using a Multimodel Ensemble. *Monthly Weather Review*, **135**(9):3239–3247 (2007).

Kynn, M. The 'heuristics and biases' bias in expert elicitation. *Journal of the Royal Statistical Society A*, **171**:239–264 (2008).

Lamarque, J.-F. Estimating the potential for methane clathrate instability in the 1%-CO_2 IPCC AR4 simulations. *Geophysical Research Letters*, **35**(19):L19806 (2008).

Landsea, C. W. & J. A. Knaff. How Much Skill Was There in Forecasting the Very Strong 1997–98 El Nino? *Bulletin of the American Meteorological Society*, **81**(9):2107–2119 (2000).

Langenfelds, R. L., *et al.* Interannual growth rate variations of atmospheric CO_2 and its $\delta^{13}C$, H_2, CH_4, and CO between 1992 and 1999 linked to biomass burning. *Global Biogeochemical Cycles*, **16**(3):21–1–21–22 (2002).

Lashof, D. A. & D. R. Ahuja. Relative contributions of greenhouse gas emissions to global warming. *Nature*, **344**(6266):529–531 (1990).

Lawrimore, J. H., *et al.* An overview of the Global Historical Climatology Network monthly mean temperature data set, version 3. *Journal of Geophysical Research*, **116**(D19):D19121 (2011).

Le Quéré, C., *et al.* Trends in the sources and sinks of carbon dioxide. *Nature Geoscience*, **2**(12):831–836 (2009).

Le Quéré, C., *et al.* Logistic constraints on CO_2 emission projections to 2050. *in prep.* (2013).

Leemans, R. & B. Eickhout. Another reason for concern: regional and global impacts on ecosystems for different levels of climate change. *Global Environmental Change*, **14**(3):219–228 (2004).

Leith, C. E. Climate response and fluctuation dissipation. *Journal of the Atmospheric Sciences*, **32**(10):2022–2026 (1975).

4507 Lenton, T. M. Land and ocean carbon cycle feedback effects on global warming in a simple
4508 Earth system model. *Tellus B*, **52**(5):1159–1188 (2000).

4509 Lenton, T. M. Early warning of climate tipping points. *Nature Climate Change*, **1**(4):201–
4510 209 (2011).

4511 Lenton, T. M., *et al.* Tipping elements in the earth's climate system. *Proceedings of the*
4512 *National Academy of Sciences*, **105**(6):1786 (2008).

4513 Leo, B. Random forests. *Machine Learning*, **45**(1):5–32 (2001).

4514 Levasseur, A., P. Lesage, & M. Margni. Considering time in LCA: Dynamic LCA and its
4515 application to global warming impact assessments. *Environmental Science & Technol-*
4516 *ogy*, **44**:2169–3174 (2010).

4517 Levitus, S., A. J., & B. T. Warming of the world ocean. *Geophysical Research Letters*, **32**
4518 (2005).

4519 Lindsay, R. W. & J. Zhang. The thinning of Arctic sea ice, 1988–2003: Have we passed a
4520 tipping point? *Journal of Climate*, **18**(22):4879–4894 (2005).

4521 Lloyd, J. & J. A. Taylor. On the temperature dependence of soil respiration. *Functional*
4522 *ecology*, **8**:315–323 (1994).

4523 Lohmann, U. & G. Lesins. Global indirect aerosol effects: A review. *Atmospheric Chem-*
4524 *istry and Physics*, **5**:715–737 (2005).

4525 Lowe, J. A., *et al.* How difficult is it to recover from dangerous levels of global warming?
4526 *Environmental Research Letters*, **4**(1):014012 (2009).

4527 Magnusson, L. & E. Källén. Factors influencing skill improvements in the ECMWF fore-
4528 casting system. *Monthly Weather Review*, **141**(9):3142–3153 (2013).

Malhi, Y., *et al.* Exploring the likelihood and mechanism of a climate-change-induced dieback of the Amazon rainforest. *Proceedings of the National Academy of Sciences of the United States of America*, **106**(49):20610–20615 (2009).

Manabe, S. & R. F. Strickler. Thermal equilibrium of the atmosphere with a convective adjustment. *Journal of the Atmospheric Sciences*, **21**(4):361–385 (1964).

Manne, A. S. & R. G. Richels. An alternative approach to establishing trade-offs among greenhouse gases. *Nature*, **410**(6829):675–677 (2001).

Matthews, H. D. Emissions targets for CO_2 stabilization as modified by carbon cycle feedbacks. *Tellus B*, **58**(5):591–602 (2006).

Matthews, H. D. & K. Caldeira. Stabilizing climate requires near-zero emissions. *Geophysical Research Letters*, **35**(4):L04705 (2008).

Matthews, H. D., S. Solomon, & R. Pierrehumbert. Cumulative carbon as a policy framework for achieving climate stabilization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **370**(1974):4365–4379 (2012).

Matthews, H. D. & A. J. Weaver. Committed climate warming. *Nature Geoscience*, **3**:142–143 (2010).

Matthews, H. D., *et al.* The proportionality of global warming to cumulative carbon emissions. *Nature*, **459**(7248):829–U3 (2009).

Mauritsen, T., *et al.* Tuning the climate of a global model. *Journal of Advances in Modeling Earth Systems*, **4**(3) (2012).

McAvaney, B., *et al. Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Model Evaluation, pages 1–54 (Cambridge University Press, 2001).

4552 McGuire, A. D., *et al.* Interactions between carbon and nitrogen dynamics in estimating net
4553 primary productivity for potential vegetation in North America. *Global Biogeochemical*
4554 *Cycles*, **6**(2):101–124 (1992).

4555 Meehl, G., *et al. Climate Change 2007: The Physical Science Basis. Contribution of Work-*
4556 *ing Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate*
4557 *Change*, chapter Global Climate Projections (Cambridge University Press, 2007).

4558 Meehl, G. A., *et al.* Combinations of natural and anthropogenic forcings in twentieth-
4559 century climate. *Journal of Climate*, **17**(19):3721–3727 (2004).

4560 Meinshausen, M., S. C. B. Raper, & T. M. L. Wigley. Emulating coupled atmosphere-ocean
4561 and carbon cycle models with a simpler model, MAGICC6–Part 1: Model description
4562 and calibration. *Atmospheric Chemistry and Physics*, **11**:1417–1456 (2011a).

4563 Meinshausen, M., T. M. L. Wigley, & S. C. B. Raper. Emulating atmosphere-ocean and
4564 carbon cycle models with a simpler model, MAGICC6–Part 2: Applications. *Atmos*
4565 *Chem Phys*, **11**(4):1457–1471 (2011b).

4566 Meinshausen, M., *et al.* Multi-gas emissions pathways to meet climate targets. *Climatic*
4567 *Change*, **75**(1-2):151–194 (2006).

4568 Meinshausen, M., *et al.* Greenhouse-gas emission targets for limiting global warming to
4569 2°C. *Nature*, **458**(7242):1158–1162 (2009).

4570 Meinshausen, M., *et al.* The RCP greenhouse gas concentrations and their extensions from
4571 1765 to 2300. *Climatic Change*, **109**(1-2):213–241 (2011c).

4572 Mignone, B. K., *et al.* Atmospheric stabilization and the timing of carbon mitigation.
4573 *Climatic Change*, **88**(3-4):251–265 (2008).

4574 Milne, G. A., *et al.* Identifying the causes of sea-level change. *Nature Geoscience*, **2**:471–
4575 478 (2009).

Mitchell, T. D. Pattern scaling: an examination of the accuracy of the technique for describing future climates. *Climatic Change*, **60**(3):217–242 (2003).

Molina, M., *et al.* Reducing abrupt climate change risk using the Montreal Protocol and other regulatory actions to complement cuts in CO_2 emissions. *Proceedings of the National Academy of Sciences*, **106**(49):20616 (2009).

Montzka, S. A., E. J. Dlugokencky, & J. H. Butler. Non-CO_2 greenhouse gases and climate change. *Nature*, **476**(7358):43–50 (2011).

Morice, C. P., *et al.* 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**(D8) (2012).

Morton, O. Climate change in black and white. *Economist* (2011).

Moss, R. H., *et al.* The next generation of scenarios for climate change research and assessment. *Nature*, **463**(7282):747–756 (2010).

Myhre, G. & D. Shindell. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Anthropogenic and Natural Radiative Forcing (Cambridge University Press, 2013).

Myhre, G., *et al.* Mitigation of short-lived heating components may lead to unwanted long-term consequences. *Atmospheric Environment*, **45**(33):6103–6106 (2011).

Nakicenovic, N., *et al. Special Report on Emissions Scenarios* (Intergovernmental Panel on Climate Change, 2000).

Narang, R. *Inside the Black Box: A Simple Guide to Quantitative and High Frequency Trading* (Wiley Finance, 2013).

Neftel, A., *et al.* Carbon Dioxide Information Analysis Center (CDIAC). *Carbon Dioxide Information Analysis Center (CDIAC)* (1994).

Nobre, C. A. & L. D. S. Borma. 'Tipping points' for the Amazon forest. *Current Opinion in Environmental Sustainability*, **1**(1):28–36 (2009).

North, G. R. & M. J. Stevens. Detecting climate signals in the surface temperature record. *Journal of Climate*, **11**(4):563–577 (1998).

O'Neill, B. C. The jury is still out on global warming potentials. *Climatic Change*, **44**(4):427–443 (2000).

Ostrom, E. How types of goods and property rights jointly affect collective action. *Journal of Theoretical Politics*, **15**(3):239–270 (2003).

Otto, A., *et al.* Climate system properties determining the social cost of carbon. *Environmental Research Letters*, **8**(2):024032 (2013a).

Otto, A., *et al.* Energy budget constraints on climate response. *Nature Geoscience*, **6**(6):415–416 (2013b).

Page, S. E., *et al.* The amount of carbon released from peat and forest fires in Indonesia during 1997. *Nature*, **420**(6911):61–65 (2002).

Pall, P., *et al.* Anthropogenic greenhouse gas contribution to flood risk in england and wales in autumn 2000. *Nature*, **470**:382–385 (2010).

Park, S. K. & R. A. Schowengerdt. Image reconstruction by parametric cubic convolution. *Computer Vision, Graphics, and Image Processing*, **23**(3):258–272 (1983).

Parry, M. L., T. R. Carter, & M. Hulme. What is a dangerous climate change? *Global Environmental Change* (1996).

Parry, M. L., J. A. Lowe, & C. Hanson. Overshoot, adapt and recover. *Nature*, **458**(7242):1102–1103 (2009).

Pawitan, Y. *In All Likelihood: Statistical Modelling and Inference Using Likelihood* (Oxford University Press, 2001).

Peck, S. C. & T. J. Teisberg. Optimal carbon emissions trajectories when damages depend on the rate or level of global warming. *Climatic Change*, **28**(3):289–314 (1994).

Peters, G. P., *et al.* Rapid growth in CO2 emissions after the 2008–2009 global financial crisis. *Nature Climate Change*, **2**(1):2–4 (2011a).

Peters, G. P., *et al.* The integrated global temperature change potential (iGTP) and relationships between emission metrics. *Environmental Research Letters*, **6**(4):044021 (2011b).

Peters, G. P., *et al.* The challenge to keep global warming below 2°C. *Nature Climate Change*, **3**(1):4–6 (2012).

Planton, S., editor. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Glossary (Cambridge University Press, 2013).

Plattner, G.-K., *et al.* Long-term climate commitments projected with climate-carbon cycle models. *Journal of Climate*, **21**(12):2721–2751 (2008).

Prentice, I., *et al.* *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, chapter The Carbon Cycle and Atmospheric Carbon Dioxide (Cambridge University Press, 2001).

Rahmstorf, S. & A. Ganopolski. Long-Term Global Warming Scenarios Computed with an Efficient Coupled Climate Model. *Climatic Change*, **43**(2):353–367 (1999).

Ramanathan, V. & D. G. Victor. To fight climate change, clear the air. *New York Times* (2010).

Ramanathan, V. & Y. Xu. The copenhagen accord for limiting global warming: Criteria, constraints, and available avenues. *PNAS*, **107**(18):8055–8062 (2010).

Ramanathan, V. R. & G. Carmichael. Global and regional climate changes due to black carbon. *Nature Geoscience*, **1**(4):221–227 (2008).

Ramaswamy, V., *et al. Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Radiative Forcing of Climate Change (Cambridge University Press, 2001).

Randall, D. A., *et al. Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Climate Models and their Evaluation (Cambridge University Press, 2007).

Raupach, M. R., J. G. Canadell, & C. L. Quéré. Anthropogenic and biophysical contributions to increasing atmospheric CO_2 growth rate and airborne fraction. *Biogeosciences*, **5**:1601–1613 (2008).

Raupach, M. R., *et al.* Global and regional drivers of accelerating CO_2 emissions. *Proceedings of the National Academy of Sciences of the United States of America*, **104**(24):10288–10293 (2007).

Raupach, M. R., *et al.* The relationship between peak warming and cumulative CO_2 emissions, and its use to quantify vulnerabilities in the carbon–climate–human system. *Tellus B*, **63**:145–164 (2011).

Rayner, S., *et al.* The Oxford Principles. *Climatic Change*, pages 1–14 (2013).

Rhein, M. & S. R. Rintoul. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Observations: Oceans (Cambridge University Press, 2013).

Rive, N., *et al.* To what extent can a long-term temperature target guide near-term climate change commitments? *Climatic Change*, **82**(3-4):373–391 (2007).

Roe, G. H. & M. B. Baker. Why Is Climate Sensitivity So Unpredictable? *Science*, **318**(5850):629–632 (2007).

Rogelj, J., D. L. McCollum, & K. Riahi. The UN's 'Sustainable Energy for All' initiative is compatible with a warming limit of 2°C. *Nature Climate Change*, **3**(6):545–551 (2013).

Rogelj, J., M. Meinshausen, & R. Knutti. Global warming under old and new scenarios using IPCC climate sensitivity range estimates. *Nature Climate Change*, **2**(4):248–253 (2012a).

Rogelj, J., *et al.* Analysis of the copenhagen accord pledges and its global climatic impacts—a snapshot of dissonant ambitions. *Environmental Research Letters*, **5**:034013 (2010).

Rogelj, J., *et al.* Discrepancies in historical emissions point to a wider 2020 gap between 2 °C benchmarks and aggregated national mitigation pledges. *Environ Res Lett*, **6**:024002 (2011a).

Rogelj, J., *et al.* Emission pathways consistent with a 2°C global temperature limit. *Nat Clim Change*, **1**(8):413–418 (2011b).

Rogelj, J., *et al.* 2020 emissions levels required to limit warming to below 2°C. *Nature Climate Change* (2012b).

Root, T., *et al.* Fingerprints of global warming on wild animals and plants. *Nature*, **421**(6918):57–60 (2003).

Ross, A. & D. Matthews. Climate engineering and the risk of rapid climate change. *Environmental Research Letters*, **4**(4):045103 (2009).

Rossby, C.-G. On the mutual adjustment of pressure and velocity distributions in certain simple current systems, ii. *Journal of Maritime Research*, **1** (1938).

Ruddiman, W. F. The anthropogenic greenhouse era began thousands of years ago. *Climatic Change*, **61**(3):261–293 (2003).

Rypdal, K., *et al.* Costs and global impacts of black carbon abatement strategies. *Tellus B*, **61B**:625–642 (2009).

Sapart, C. J., *et al.* Natural and anthropogenic variations in methane sources during the past two millennia. *Nature*, **490**(7418):85–88 (2013).

Sausen, R. & U. Schumann. Estimates of the climate response to aircraft CO_2 and NO_x emissions scenarios. *Climatic Change*, **44**:27–58 (2000).

Schlesinger, M. E. & X. Jiang. Simple model representation of atmosphere-ocean GCMs and estimation of the time scale of CO_2-induced climate change. *Journal of Climate*, **3**:1297–1315 (1990).

Schneider, S. Abrupt non-linear climate change, irreversibility and surprise. *Global Environ Chang*, **14**(3):245–258 (2004).

Schneider, S. H. What is 'dangerous' climate change? *Nature*, **411**(6833):17–19 (2001).

Sellers, W. D. A global climatic model based on the energy balance of the earth-atmosphere system. *Journal of Applied Meteorology*, **8**:392–400 (1969).

Shindell, D., *et al.* Simultaneously mitigating near-term climate change and improving human health and food security. *Science*, **335**(6065):183–189 (2012).

Shine, K. The global warming potential — the need for an interdisiplinary retrial. *Climatic Change*, **96**:467–472 (2009).

Shine, K. P., *et al.* Alternatives to the global warming potential for comparing climate impacts of emissions of greenhouse gases. *Climatic Change*, **68**(3):281–302 (2005).

Shine, K. P., *et al.* Comparing the climate effect of emissions of short- and long-lived climate agents. *Philosophical Transactions of the Royal Society A*, **365**(1856):1903–1914 (2007).

Siegenthaler, U. & F. Joos. Use of a simple model for studying oceanic tracer distributions and the global carbon cycle. *Tellus B*, **44**(3):186–207 (1992).

Simmons, A. J. & A. Hollingsworth. Some aspects of the improvement in skill of numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, **128**(580):647–677 (2002).

Smith, J. B., *et al.* Assessing dangerous climate change through an update of the Intergovernmental Panel on Climate Change (IPCC) "reasons for concern". *Proceedings of the National Academy of Sciences of the United States of America*, **106**(11):4133–4137 (2009).

Smith, S. M., *et al.* Equivalence of greenhouse-gas emissions for peak temperature limits. *Nat Clim Change*, **2**(5):1–4 (2012).

Sokolov, A. P. & P. H. Stone. A flexible climate model for use in integrated assessments. *Climate Dynamics*, **14**(4):291–303 (1998).

Solomon, S., *et al.* Irreversible climate change due to carbon dioxide emissions. *Proceedings of the National Academy of Sciences*, **106**(6):1704–1709 (2009).

4736 Sotillo, M. G., *et al.* A high-resolution 44-year atmospheric hindcast for the Mediterranean
4737 Basin: contribution to the regional improvement of global reanalysis. *Climate Dynamics*,
4738 **25**(2-3):219–236 (2005).

4739 Stainforth, D., *et al.* Uncertainty in predictions of the climate response to rising levels of
4740 greenhouse gases. *Nature*, **433**(7024):403–406 (2005).

4741 Stefanova, L., *et al.* Hindcast skill and predictability for precipitation and two-meter air
4742 temperature anomalies in global circulation models over the Southeast United States.
4743 *Climate Dynamics*, **38**(1-2):161–173 (2011).

4744 Steinacher, M., F. Joos, & T. F. Stocker. Allowable carbon emissions lowered by multiple
4745 climate targets. *Nature*, **499**(7457):197–201 (2013).

4746 Stern, D. I. & R. K. Kaufmann. Estimates of global anthropogenic methane emissions
4747 1860–1993. *Chemosphere*, **33**(1):159–176 (1996).

4748 Stocker, T. *Introduction to Climate Modelling* (Springer, 2011).

4749 Stone, D. A. Attribution of global surface warming without dynamical models. *Geophysi-*
4750 *cal Research Letters*, **32**(18):L18711 (2005).

4751 Stott, P. A. & J. A. Kettleborough. Origins and estimates of uncertainty in predictions of
4752 twenty-first century temperature rise. *Nature*, **416**:723–726 (2002).

4753 Stott, R. Healthy response to climate change. *BMJ*, **332**(7554):1385–1387 (2006).

4754 Stroeve, J. C., *et al.* Trends in Arctic sea ice extent from CMIP5, CMIP3 and observations.
4755 *Geophysical Research Letters*, **39**(16) (2012).

4756 Tarnocai, C., *et al.* Soil organic carbon pools in the northern circumpolar permafrost region.
4757 *Global Biogeochemical Cycles*, **23**(2) (2009).

4758 Taylor, F. W. *Elementary Climate Physics* (Oxford University Press, 2005).

4759 Taylor, G. I. Eddy motion in the atmosphere. *Philosophical Transactions of the Royal*
4760 *Society A*, **215**:1–26 (1915).

4761 Taylor, K. E., R. J. Stouffer, & G. A. Meehl. An Overview of CMIP5 and the Experiment
4762 Design. *Bulletin of the American Meteorological Society*, **93**(4):485–498 (2012).

4763 Tebaldi, C. & R. Knutti. The use of the multi-model ensemble in probabilistic climate
4764 projections. *Philosophical Transactions of the Royal Society A: Mathematical, Physical*
4765 *and Engineering Sciences*, **365**(1857):2053–2075 (2007).

4766 Teixeira, J., C. A. Reynolds, & K. Judd. Time Step Sensitivity of Nonlinear Atmospheric
4767 Models: Numerical Convergence, Truncation Error Growth, and Ensemble Design. *Jour-*
4768 *nal of the Atmospheric Sciences*, **64**(1):175–189 (2007).

4769 Thompson, D. W. J., *et al.* A large discontinuity in the mid-twentieth century in observed
4770 global-mean surface temperature. *Nature*, **453**(7195):646–649 (2008).

4771 Tickell, O. *Kyoto2* (Zed Books Ltd, 2008).

4772 Tol, R. S. J. The marginal costs of greenhouse gas emissions. *The Energy Journal*,
4773 **20**(1):61–81 (1999).

4774 Tol, R. S. J., *et al.* A unifying framework for metrics for aggregating the climate effect of
4775 different emissions. *Environmental Research Letters*, **7**(4):044006 (2012).

4776 Trenberth, K. E. & P. D. Jones. *Climate Change 2007: The Physical Science Basis. Con-*
4777 *tribution of Working Group I to the Fourth Assessment Report of the Intergovernmen-*
4778 *tal Panel on Climate Change*, chapter Observations: Surface and Atmospheric Climate
4779 Change (Cambridge University Press, 2007).

4780 UK Committee on Climate Change. *Building a low-carbon economy - the UK's contribu-*
4781 *tion to tackling climate change* (The Stationary Office, 2008).

UK Government. Reducing the UK's greenhouse gas emissions by 80(2012a).

UK Government. Supporting international action on climate change (2012b).

UNEP. *Emissions Gap Report* (United National Environment Programme (UNEP), Nairobi, 2012).

UNEP & WMO. *Integrated Assessment of Black Carbon and Tropospheric Ozone: Summary for Decision Makers* (United Nations Environment Programme (UNEP), Nairobi, 2011).

UNFCCC. Copenhagen Accord (2009).

UNFCCC. Report of the Conference of the Parties on its sixteenth session, held in Cancun from 29 November to 10 December 2010 (2010).

UNFCCC. UNFCCC establishment of an Ad Hoc Working Group on the Durban Platform for Enhanced Action (2011).

United Nations. United Nations Framework Convention on Climate Change (1992).

United Nations. Kyoto Protocol to the United Nations Framework Convention on Climate Change (1998).

Updegraff, K., *et al.* Hysteresis in the temperature response of carbon dioxide and methane production in peat soils. *Biogeochemistry*, **43**(3):253–272 (1998).

Vallis, G. K. *Atmospheric and Oceanic Fluid Dynamics* (Cambridge University Press, 2006).

van den Broeke, M., *et al.* Temperature thresholds for degree-day modelling of Greenland ice sheet melt rates. *Geophysical Research Letters*, **37**(18) (2010).

van Vuuren, D. P., *et al.* The representative concentration pathways: an overview. *Climatic Change*, **109**(1-2):5–31 (2011).

Vaughan, N. E. & T. M. Lenton. A review of climate geoengineering proposals. *Climatic Change*, **109**(3-4):745–790 (2011).

Vaughan, N. E. & T. M. Lenton. Interactions between reducing CO_2 emissions, CO_2 removal and solar radiation management. *Philosophical Transactions of the Royal Society A*, **370**(1974):4343–4364 (2012).

Vaughan, N. E., T. M. Lenton, & J. G. Shepherd. Climate change mitigation: trade-offs between delay and strength of action required. *Climatic Change*, **96**(1-2):29–43 (2009).

Vellinga, M. & R. A. Wood. Global climatic impacts of a collapse of the Atlantic thermohaline circulation. *Climatic Change*, **54**(3):251–267 (2002).

Victor, D. G., C. F. Kennel, & V. Ramanathan. The climate threat we can beat. *Foreign Affairs*, **91**:112–114 (2012).

Vuuren, D. P., *et al.* How well do integrated assessment models simulate climate change? *Climatic Change*, **104**(2):255–285 (2009).

Vuuren, D. P., *et al.* Rcp2.6: exploring the possibility to keep global mean temperature increase below 2°C. *Climatic Change*, **109**(1-2):95–116 (2011).

Wallack, J. S. & V. Ramanathan. The other climate changers: Why black carbon and ozone also matter. *Foreign Affairs*, **88**:105–113 (2009).

Wang, W.-C., M. Dudek, & X.-Z. Liang. Inadequacy of effective CO_2 as a proxy in assessing the regional climate change due to other radiatively active gases. *Geophysical Research Letters*, **19**:1375–1378 (1992).

Wang, W.-C., *et al.* Inadequacy of effective CO_2 as a proxy in simulating the greenhouse effect of other radiatively active gases. *Nature*, **350**:573–577 (1991).

WBGU. Solving the climate dilemma: the budget approach (2009).

Weaver, A. J., *et al.* Long term climate implications of 2050 emission reduction targets. *Geophysical Research Letters*, **34**(19):L19703 (2007).

Wendisch, M. & P. Yang. *Theory of Atmospheric Radiative Transfer: A Comprehensive Introduction* (Wiley, 2012).

Wigley, T. M. L. A Combined Mitigation/Geoengineering Approach to Climate Stabilization. *Science*, **314**(5798):452–454 (2006).

Williams, R. G. & M. J. Follows. *Ocean Dynamics and the Carbon Cycle: Principles and Mechanisms* (Cambridge University Press, 2011).

Winton, M. Do climate models underestimate the sensitivity of Northern Hemisphere sea ice cover? *Journal of Climate*, **24**(15):3924–3934 (2011).

Wu, P., *et al.* Temporary acceleration of the hydrological cycle in response to a CO_2 rampdown. *Geophysical Research Letters*, **37**(12) (2010).

Zarzycki, C. M. & T. C. Bond. How much can the vertical distribution of black carbon affect its global direct radiative forcing? *Geophysical Research Letters*, **37**(20) (2010).

Zickfeld, K., V. K. Arora, & N. P. Gillett. Is the climate response to CO_2 emissions path dependent? *Geophysical Research Letters*, **39**(5) (2012).

Zickfeld, K., *et al.* Setting cumulative emissions targets to reduce the risk of dangerous climate change. *Proceedings of the National Academy of Sciences*, **106**(38):16129–16134 (2009).

Zickfeld, K., *et al.* Expert judgments about transient climate response to alternative future trajectories of radiative forcing. *Proceedings of the National Academy of Sciences* (2010).

Zickfeld, K., *et al.* Long-Term Climate Change Commitment and Reversibility: An EMIC Intercomparison. *Journal of Climate*, **26**(16):5782–5809 (2013).