Characterizing Impacts and Implications of Proposals for Solar Radiation Management, a Form of Climate Engineering

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in
the Department of Engineering and Public Policy

Katharine L. Ricke

S.B., Earth, Atmospheric and Planetary Sciences,
Massachusetts Institute of Technology

Carnegie Mellon University
Pittsburgh, Pennsylvania

August, 2011

Abstract

Even under optimistic emissions scenarios, rising concentrations of greenhouse gases in the atmosphere will result in significant increases in global mean temperatures and associated effects for the foreseeable future (IPCC, 2007a,b). Concerns that mitigation may be too slow in coming have lead to renewed dialogue within the scientific community regarding potential strategies for counteracting global warming through geoengineering, defined as "the deliberate large-scale intervention in the Earth's climate system, in order to moderate global warming." (Shepherd et al., 2009)

The geoengineering schemes that are considered most feasible today involve planetary albedo modification, or "solar radiation management" (SRM). This thesis addresses several outstanding questions regarding uncertainty in global and regional effects of SRM activities. The technical components of this work are centered on two modeling experiments which use a coupled atmosphere-ocean general circulation model (AOGCM) implemented through climate*prediction*.net. Drawing upon knowledge gained through these experiments and interaction with the broader research community, I explore the international relations implications of SRM and the global governance issues associated with it.

The first experiment explored regional differences in climate modified by SRM using a large-ensemble modeling experiment that examines the effects of 54 global temperature stabilization scenarios. Our results confirm other research that shows a world with SRM would generally have less extreme temperature and precipitation anomalies than one with unmitigated greenhouse gas emissions and no SRM, but illustrate the physical unfeasibility of simultaneously stabilizing global precipitation and temperature as long as greenhouse gases continue to rise. Over time, simulated temperature and precipitation in large regions such as China and India vary significantly with different SRM trajectories and diverge from historic baselines in different ways. Hence the use of SRM to stabilize climate in all regions simultaneously

may not be possible. Regional diversity in the response to different levels of SRM could complicate what is already a very challenging problem of global governance, and could make consensus about the "optimal" level of geo-engineering difficult, if not impossible, to achieve.

The second experiment modeled SRM using a perturbed physics ensemble with a wide range of temperature responses and climate sensitivities, all of which are consistent with observed recent warming. The analysis shows that the efficacy and distribution of effects of SRM varies with the temperature response of the model. Models that produce more global warming are also generally more sensitive to SRM, so the amount of modification of the Earth's energy balance needed to meet any given climate stabilization criteria appear to be relatively insensitive to climate sensitivity. While in the more sensitive models, SRM is generally less successful in returning regional climates to their unperturbed states the longer it is used to compensate for rising greenhouse gases, it is also where SRM is most effective relative to a no SRM alternative.

SRM does not prevent further acidification of the oceans and this fact, coupled with the fact that SRM can only slow, never halt, changes to regional climate states, makes SRM untenable as a long-term solution to the problems caused by rising GHGs in the atmosphere. Much more research on SRM is needed before any conclusions on whether or not to deploy it are reached, but this work suggests that regional inequities in climate response are probably not the main impediment to its effective implementation. While SRM can never perfectly correct for regional climate change, these experiments suggest that it generally reduces (rather than exacerbates) changes to regional temperature and precipitation and greatly reduces the rate of temperature change. Considering the slow progress society has made towards reducing emissions, however, it is important to consider the potential benefits SRM technologies may confer in reducing impacts to buy time for both mitigation and adaptation.

Acknowledgments

My thesis proposal last year was one of the most gratifying experiences of my professional life thus far. It was an honor to have four such distinguished scientists debating my work, and I would like to thank my committee for their excellent guidance. I have benefitted greatly from having M. Granger Morgan and David W. Keith as my own personal red team-blue team of advisors and committee co-chairs. By pushing me to always "think more carefully" about my assumptions and my conclusions, David's skepticism has forced me to make my work better and more honest. Granger has been an amazing enabler, a wonderful teacher, mentor and advocate for me and my work. I have been incredibly humbled by his generosity with his time and wisdom, and feel so lucky to have gotten the chance to be his student. Jay Apt has offered me a host of guidance over the past four years on my research and professional development and I appreciate all his insightful comments. Myles Allen, I cannot thank enough for hosting me in his group and giving me the great opportunity to work with climate*prediction*.net. Had I not been lucky enough to be invited to dinner with him after his seminar in the spring of 2008, this thesis would have looked completely different.

In addition to Myles, it has been my privilege to work with a number of other great scientists at Oxford. I am especially indebted to Dan Rowlands for collaborating with me on the ensemble design for the second climate*prediction*.net geoengineering experiment. I'd like to thank Philip Stier, Hiro Yamazaki, Neil Massey, Kuniko Yamazaki and Sue Rosier for the advice and guidance in the summer of 2008 as I was setting up my first experiment; Tolu Aina, Milo Thurston and Neil again for making my experiments go and helping me access the data; and William Ingram for invaluable assistance obtaining information on the HadCM3 and extensive, frank feedback on the analysis and interpretation of results from both experiments. In addition, this work would not have been possible without the

generous donations of personal computer power from thousands of climate*prediction*.net participants.

In EPP, I have benefitted greatly from interactions with many faculty members, in particular Benoït Morel and Peter Adams. The staff in EPP are also amazing, and Patti, Meryl, Barb, Vicki, Julie and Adam have all provided me with a ton of assistance. A number of other researchers have been extremely supportive of my work and provided invaluable feedback on it over the past four years, including Klaus Keller, Ken Caldeira, Jane Long, Alan Robock, Doug MacMynowski, Iris Grossmann and especially Juan Moreno Cruz, a great collaborator and friend.

I'd like to thank my friends and fellow students here at CMU, who I have learned much from and relied on for support and assistance in tough times, especially Anu, Sharon, Lauren, Jon, Tim, Royce, and many others! Thanks also to my MIT friends, especially Teri, Victor, Dan, Marc, Gretchen and Andrew who have continued to make life awesome even as we have spread out to grad schools and other endeavors coast-to-coast. Thank you to Smidgen, my best furry friend, RIP.

I have been lucky to have a number of amazing mentors who have shaped my career path greatly. I want to thank my high school science teachers, Jim Tuchscherer and Hank Ryan; my first real research supervisor, Professor John Dickey, formerly of the Physics/Astronomy Department at the University of MN; my MIT academic advisor, UROP supervisor and undergraduate thesis advisor, Professor Julian Sachs, now at University of Washington; Susan Murcott in the Environmental Engineering Department at MIT, who supported me through some of the most formative experience of my life while doing field research in Nicaragua and Thailand, and the late, wonderful Ron Rivera, who convinced me to focus on learning for a bit before trying to save the world; and Andrew Stoeckle and Susan Day at Abt Associates.

Thank you to my family for their support, including my two sisters, Margaret and

Isabel, for being the coolest; and my parents, for giving me all the opportunities that allowed me to get to this point. I am grateful to my mom for pushing me and providing me with a great role model for a strong female professional in a male-dominated field. My dad has been my biggest support system for as long as I can remember. Thank you, Dad, for all you have done for me and helped me achieve.

Finally, I would like to thank the best husband ever, Marcos Ojeda, for support that went far beyond the call of duty for a graduate student spouse. Thank you for making me switch to Python and your many hours of help with my code! Thank you for all your help with Adobe Illustrator! Thank you for proof reading my thesis and figuring out how to use this nice font in LaTeX! Thank you for making me food, bringing me coffee, picking me up at school at all hours and helping me realize my dream of wheeling down Baker Hall in my desk chair without running into the wall. You are the absolute best!

This research was made possible through support from a National Science Foundation Graduate Research Fellowship; the Climate Decision Making Center (CDMC), created through a cooperative agreement between the National Science Foundation (SES-0345798) and Carnegie Mellon University; the Center for Climate and Energy Decision Making (CEDM), created through a cooperative agreement between the National Science Foundation (SES-0949710) and Carnegie Mellon University; and the Pittsburgh branch of the ARCS Foundation, and in particular, Beth Wainwright and Francine Abraham, the donors of my named ARCS scholarship.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

According to the IPCC Fourth Assessment, even under optimistic emissions scenarios, rising concentrations of greenhouse gases in the atmosphere will result in significant increases in global mean temperatures and associated impacts for the foreseeable future. (IPCC, 2007a,b) Concerns that mitigation may be too slow in coming have lead to renewed dialogue within the scientific community regarding potential strategies for counteracting global warming through geoengineering.

Geoengineering is defined as "the deliberate large-scale intervention in the Earth's climate system, in order to moderate global warming." (Shepherd et al., 2009) It differs from abatement or mitigation in that not all geoengineering schemes would reduce the greenhouse gas concentrations in the atmosphere that are the root cause of climate change. It also differs from adaptation in that geoengineering aims to modify the climate system directly to ameliorate the effects of global warming, rather than make adjustments to human systems to reduce damages or suffering.

Cooling down the climate using geoengineering is not a new idea - in fact, a 1965 report to US President Johnson cited geoengineering by increasing planetary albedo as the sole proposed solution to the emerging problem of rising atmospheric CO_2. (Keith, 2000a)

Only recently, however, has the idea gained a broad degree of legitimacy as several prominent climate scientists have raised it as a feasible, and potentially necessary, strategy for avoiding catastrophic impacts of climate change.

In particular, a 2006 article in *Climatic Change* by Nobel laureate Paul Crutzen proposed injecting sulphates into the stratosphere to counteract the effect of removing them from the troposphere (through pollution control measures). (Crutzen, 2006) Soon after, UCAR senior scientist, Tom Wigley, wrote a piece in *Science* that used a simple modeling exercise to examine the possibility of stabilizing global temperatures with a combination of mitigation and geoengineering. (Wigley, 2006) Since then, a multitude of studies by climate scientists have been published on the topic, and several coordinated research projects have been initiated to systematically further scientific understanding of the subject.

The geoengineering schemes that are considered most feasible today involve planetary albedo modification. Albedo is the fraction of incoming solar radiation that is reflected back into space, rather than being absorbed by the earth and its atmosphere. If this reflectivity is increased, the amount of the Sun's energy absorbed by the Earth will decrease and global temperatures will drop. Activities that attempt to exploit this mechanism are called "solar radiation management," or SRM, for the remainder of this document.

The thesis addresses several outstanding questions regarding uncertainty in global and regional effects of SRM activities. The technical components of this work are centered on two modeling experiments which use a coupled atmosphere-ocean general circulation model (AOGCM) implemented through climate*prediction*.net. Drawing upon knowledge gained through these experiments and interaction with the broader research community, I explore the international relations implications of SRM and the global governance issues associated with it.

1.1 Why Consider Managing Solar Radiation?

The motivation behind studying SRM bears some similarity to that for studying climate change adaptation. Both address the fact of some inevitable warming. While adaptation can be thought of as a way to address moderate and expected changes to the climate by modifying human behavior and infrastructure, SRM can be thought of as a way to address extreme changes to the climate through direct intervention. (Victor, 2011)

Figure 1.1 shows projected global temperature changes for different anthropogenic emissions scenarios as presented in the IPCC Fourth Assessment Report (Meehl et al., 2007). Even under optimistic emissions scenarios, such as SRES B1– which is predicated on low population growth and low energy use – nearly 2°C of warming from preindustrial is expected by the end of the century. The real uncertainty associated with anthropogenic perturbations to the climate system is even greater than shown in Figure 1.1. Standard versions of general circulation models (GCMs) cluster around the median of the distribution of climate sensitivity, so the observationally constrained distributions in this figure are skewed, and well below its substantial high tail (Roe and Baker, 2007).

Uncertainty about the amount of warming expected in the future is accompanied by uncertainty about the amount of warming required to reach certain "tipping points" that, if passed, would lead to fundamental and perhaps irreversible shifts in the way the climate system functions. Such tipping elements include the disappearance of Arctic summer sea-ice, collapse of the ocean's thermohaline circulation or rapid decay of the Greenland ice sheet. (Lenton et al., 2008) These are the types of climate changes, in particular, that could produce rapid or "catastrophic" impacts that would be difficult to adapt to. While the probabilities associated with climate catastrophes may be low, their dire economic consequences contribute disproportionately to the risks associated with global warming (Weitzman, 2009). The extent to which SRM could stop or slow the progression of extreme

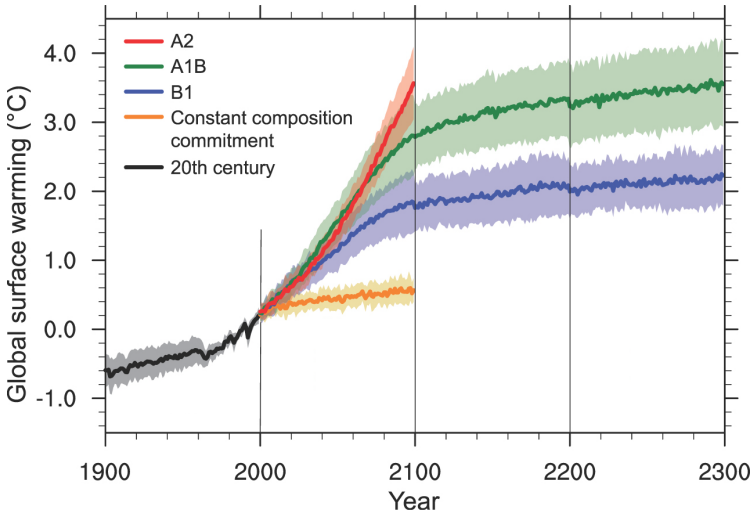

Figure 1.1: Multi-model means of surface warming (relative to 1980–1999) for the scenarios A2, A1B and B1, shown as continuations of the 20th-century simulation. (Figure 10.4 from Chapter 10 of the the IPCC Fourth Assessment Report of Working Group 1. (Meehl et al., 2007))

climate change after its commencement or detection is still unknown, but seeking insurance against the worst-case scenarios is the primary reason cited for investing in SRM research. (Blackstock et al., 2009; Morgan and Ricke, 2010; Shepherd et al., 2009)

1.2 Characteristics of SRM

A number of different approaches have been proposed to deflect more incoming solar radiation before it is absorbed by the Earth. These include surface albedo enhancements, such as crop modifications or desert reflectors; marine cloud albedo enhancement through cloud seeding; stratospheric albedo modification with aerosols; and space-based reflectors, either in low Earth orbit or at the Lagrange L1 point. Figure 1.2 from Shepherd et al. (2009) shows how these different placements of SRM technologies could modify the Earth's radiative fluxes.

Stratospheric SRM is the focus of this thesis and the majority of work on SRM thus far

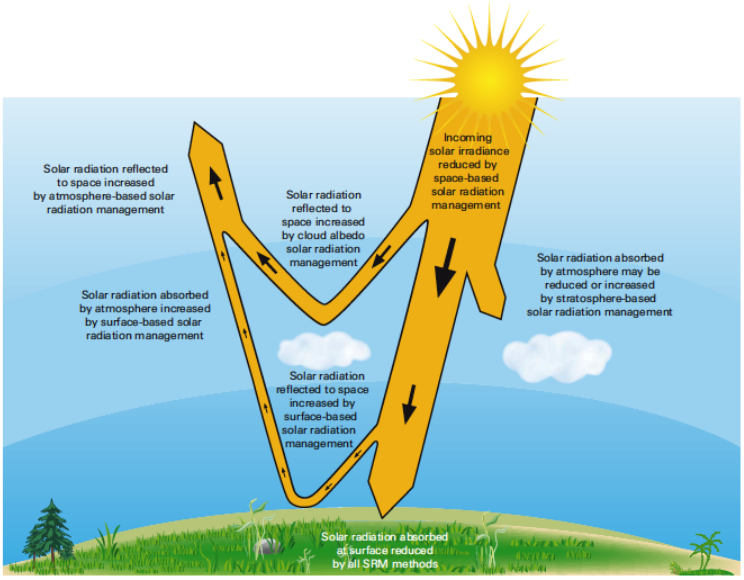

Figure 1.2: Schematic showing the impact of different SRM methods on solar radiation fluxes. (Figure 3-1 from the Royal Society's "Geoengineering the Climate: Science, governance and uncertainty". (Shepherd et al., 2009))

because of its well understood natural analog (volcanoes), its global scalability (Lenton and Vaughan, 2009) and some substantive technical understanding of how it could be implemented. (McClellan et al., 2010) Whatever the approach, however, SRM can be characterized by "three essential characteristics: it is cheap, fast and imperfect." (Keith et al., 2010)

CHEAP

The classification of SRM activities as "cheap" doesn't just refer to the low economic costs associated with cooling the planet with these mechanisms, but also to the fact that only a little bit of material is necessary to implement these planetary-scale changes, which can offset the influence of tons of CO_2. For example, under the current understanding of SRM technologies, the mass of fine particles needed to counteract the radiative effects of a doubling of atmospheric CO_2 concentrations is approximately 2.6 million tons per day of

aerosol if injected into marine stratus clouds or 13 thousands tonnes per day of sulphate aerosol if injected into the stratosphere. By comparison, to achieve the same radiative effect (whether by artificial or natural means), we would need to remove 225 million tons per day of CO2 from the atmosphere for 25 years straight. (Keith, 2009)

While few realistic engineering analyses have been done on the economic costs of SRM, a 1992 report of the National Research Council estimated the potential costs of a program of stratospheric albedo modification based on the use of a standard naval gun system dispensing commercial aluminium oxide dust to counteract the warming effect of a CO2 doubling. Undiscounted annual costs for a 40-year project were estimated to be $100-billion(NAS, 1992). More recent analyses (McClellan et al., 2010; Robock et al., 2009; Salter et al., 2008), have suggested that well designed systems might reduce this cost to less than $10-billion per year – clearly well within the budget of most countries, and much less costly than any program to dramatically reduce the emissions of CO_2.

FAST

While cutting emissions of CO_2 and other greenhouse gases would slow or halt their rising concentrations in the atmosphere, much of the CO_2 released through past emissions will reside in the atmosphere for 100 years or more. In addition, inertia in the climate system means that global temperatures will continue to rise. Reducing planetary temperatures through emissions reductions will take many decades to centuries. In contrast, increasing planetary albedo via SRM can reduce planetary temperature in days or months. This fast response cuts two ways. On the one hand, it means that SRM could be used to rapidly cool the planet in the event of a "climate emergency", such as the rapid deterioration of the Greenland ice sheet or the sudden release of large amounts of methane from arctic tundra or the deep edges of the coastal oceans. On the other hand, if SRM were started and

then stopped before greenhouse gas concentrations in the atmosphere were drastically reduced, then global temperatures could shoot up dramatically (Matthews and Caldeira, 2007). This would be devastating for many ecosystems.

IMPERFECT

Because the mechanisms by which SRM cools the planet are different from those by which greenhouse gases warm it, SRM cannot reverse climate change in a perfect way at either the global or local level. Global warming from rising greenhouse gases changes the level of global precipitation in a number of ways. First, rising global temperatures cause more evaporation, which in turn, produces more precipitation. Second, higher concentrations of greenhouse gases in the atmosphere also reduces the amount of radiative cooling of the troposphere at any given temperature, which must be balanced by a reduction in latent heating (therefore precipitation). (Yang et al., 2003) With rising concentrations of greenhouse gases and rising global temperatures, the precipitation-increasing effect dominates and precipitation increases globally. But, when SRM is used to lower global temperatures in a world with high CO_2, only the second dampening effect remains. Thus, SRM necessarily weakens the global hydrological cycle. (Bala et al., 2008) This effect would affect different regions of the planet differently, with SRM compensating for climate changes in some regions reasonably well but potentially exacerbating climate change-driven effects in others. It is almost certain that the benefits and costs of global climate stabilization would not be equitably distributed among regions. (Ricke et al., 2010; Robock et al., 2008)

In addition to such imperfections, a number of negative side effects could result from the various proposals for implementing SRM. Injecting aerosols into the stratosphere could provide reaction sites that might lead to significant destruction of stratospheric

ozone. (Tilmes et al., 2009) And because SRM does nothing to stop the rise of CO_2 from anthropogenic activity, it cannot slow the associated acidification of the surface ocean, the continuation of which could lead to profound changes in ocean and terrestrial ecosystems, including the likely demise of many or all coral reefs. (Doney et al., 2009)

1.3 Mechanisms of SRM

Volcanic Eruptions as an Analog for Stratospheric SRM

Stratospheric albedo modification, achieved by increasing the amount of dust or aerosols in the stratosphere, is the SRM strategy that has received the most attention and research time thus far. The climatic dynamics of this process are fairly well understood because of an analogous natural phenomenon—large, explosive volcanic eruptions. When such a volcanic eruption occurs, sulfur dioxide and some hydrogen sulfide are blasted into the stratosphere, where they are converted to sulphate aerosols resulting in observed decreases in global temperature. (Robock, 2003) For example, the eruption of Mount Pinatubo in the Philippines in 1991 produced global scale cooling of about 0.5 $^\circ$ C. (Soden et al., 2002)

The planet-cooling processes of volcanic eruptions have two properties attractive to geoengineers. First, changes in temperature begin to occur immediately after an eruption, meaning such mechanisms could be used for a rapid response to a climate catastrophe. In addition, because the stratosphere is convectively stable, particles can reside there for one to two years (as opposed to just a few days in the unstable troposphere).(Robock, 2003) Thus, the system requires replacement on only an annual or perhaps biannual basis, as opposed to constant replenishment, but still allows for tuning or phasing out over time.

While the stratospheric aerosols injected by volcanic eruptions produce a net global

cooling, they also cause warming of the stratosphere and winter warming of the troposphere over the continental Northern Hemisphere, which can in turn change mean atmospheric circulation. (Robock, 2003) Both the El Chichon eruption in 1982 and Mount Pinatubo eruption in 1991 were associated with decreases in precipitation over land, and there were large decreases in surface runoff after Pinatubo. (Trenberth and Dai, 2007) Modelling studies and paleoclimate data indicate that high latitude volcanic eruptions disrupt the African and Indian monsoon cycles. (Oman et al., 2006)

SRM and the Hydrological Cycle

As mentioned above, any SRM scheme that aims to stabilize tropospheric temperatures by canceling the longwave forcing of increased tropospheric CO_2 with a shortwave forcing of reduced insolation will physically reduce the intensity of the hydrological cycle. As illustrated in Figure 1.3a, absent some external forcing perturbation, to maintain a constant global-mean temperature in the troposphere its radiative cooling must be balanced by latent heating of condensation of water vapor (Yang et al., 2003). When concentrations of CO_2 and other GHGs in the troposphere increase, temperatures and specific humidity will rise according to the Clausius-Clayperon relation which results in an increase in precipitation. Independent of temperature changes, however, rising concentrations of CO_2 will decrease radiative cooling ability of the troposphere. This will be balanced by a decrease in its ability to release latent heat (thus a decrease in precipitation) (Allen and Ingram, 2002). Andrews et al. (2009) characterizes this second effect as the "fast" response of precipitation to CO_2, as it is observed in simulations almost immediately after a CO_2 perturbation is made (Fig. 1.3b). Eventually, as the surface and troposphere warm, both radiative cooling and latent heating (precipitation) rise proportionately, according to Clausius-Clapeyron (Fig. 1.3c). Most models show this "slow" response of precipitation dominates the "fast,"

which is why global-mean precipitation is expected (and has been observed) to increase under global warming.

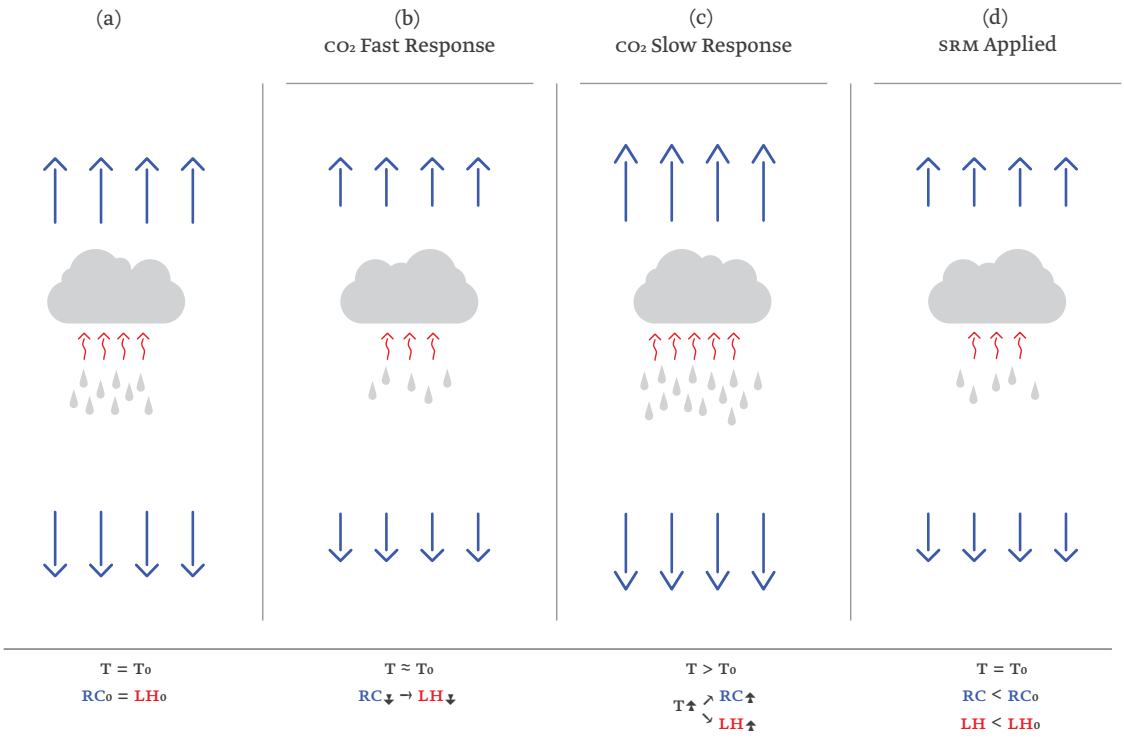

Figure 1.3: Schematic showing the effects of CO_2 (b-c) and SRM (d) on radiative cooling, latent heating and temperature of the troposphere.)

When a SRM is used to constrain tropospheric temperatures, the "slow" effect of CO_2 on precipitation (i.e., temperature dependent) effect is essentially eliminated and what remains is the "fast" precipitation dampening effect of CO_2 (Fig. 1.3d). To the first-order, this explains why physically global-mean precipitation must fall in an SRM-stabilized troposphere with rising CO_2 concentrations. This explanation and schematic represent a simplified description of the effects of SRM on the tropospheric energy budget and the global-mean hydrological cycle. In reality, feedbacks result in small changes to the shortwave radiation and sensible heat fluxes in the troposphere in addition to the longwave radiative cooling and latent heating changes. (See Andrews et al., 2009.)

Properties and Limitations of Proposed Stratospheric Particles

A number of potential types of particles or particle precursors have been suggested for injection to increase the albedo of the stratosphere, including sulfur dioxide, aluminum oxide dust or even designer self-levitating aerosols that could be engineered to migrate to particular regions (e.g. over the arctic) or to rise above the stratosphere (so as not to interfere in stratospheric chemistry). The most desirable particles would provide maximum scattering efficiency, minimal infrared absorption and low (chemical and photo-) reactivity using a small amount of mass.(Blackstock et al., 2009)

Teller et al. (1997) and Teller et al. (2002) explored in some detail different types of particles, contrasting their efficiencies and costs. Particles similar to those produced by volcanic eruptions (dielectric oxides) would be less efficient scatterers than conductive metals which, in turn would be less efficient than hypothetical engineered "quasi-resonant scatterers" that exploit atomic transition-frequency of the particle material. Keith (2010) explored the possibility of engineering self-levitating particle that exploit photophoretic forces and could increase flexibility of placement of scatterers in the atmosphere and have longer lifetimes and fewer side effects than sulphates.

Because of their volcanic analog, most practical discussions of creating sulphate aerosols have suggested using SO_2 as a precursor gas. Recent literature, however, suggests that there may be some physical limitations to the effectiveness of this approach and other precursor gases or particle may be more desirable. Heckendorn et al. (2009) shows that as the injection of SO_2 increases, the corresponding sulphate aerosol burden-to-injection ratio decreases as does the amount of radiative forcing-to-injection ratio. Pierce et al. (2010) showed that better aerosol size distributions could be achieved by injecting H_2SO_4 directly. As mentioned above, aerosols injected into the stratosphere could lead to significant destruction of stratospheric ozone (Tilmes et al., 2009). They could also lead to

increased stratospheric ozone through temperature feedbacks (Crutzen, 2006).

Cloud Albedo Modificaton

While the focus of this thesis is stratospheric SRM, its effects share some significant similarities, as well as some relative advantages and disadvantages, with proposals to enhance planetary albedo through marine cloud modification. The basic idea behind marine cloud albedo modification is to enhance the reflectivity of certain types of clouds by exploiting the Twomey effect, by which the increase of cloud condensation nuclei (CCN) increases the number of droplets in a cloud, decreases their particle size and thereby increases their brightness. (Twomey, 1977) There is an observable analog associated with this type of SRM, as well — the tracks observed from exhaust released from ships. Salter et al. (2008) proposes a design for ships whose express purpose would be to enhance marine cloud brightness. Latham et al. (2008) demonstrated that such an approach could work in principle, but it is still unclear whether it could implemented over a large enough area of the oceans to have the significant negative forcing effect needed to counter anthropogenic climate change.

Because, like stratospheric SRM, cloud albedo modification compensates for increasing anthropogenic longwave forcings with enhanced shortwave forcings, the approach has some similar effects. Some results exploring these climate effects of marine cloud enhancement are presented, and contrasted with those of stratospheric SRM, in the following section on modeling of SRM.

1.4 Modeling of SRM

An ever expanding number of climate modeling exercises to examine the climatic effects of SRM have been published in the past decade. As the primary focus of this document

is the results from two climate modeling experiments, the following section presents an overview of key results from these studies.

Scoping Exercises with Lower Complexity Models

Experiments using basic energy balance or limited complexity models have been able to reveal certain dynamics of SRM at the global scale that have guided some elements of the experimental design in subsequent experiments with more sophisticated models. Wigley (2006) used a simple energy balance to show how SRM could be used to buy time for mitigation, stabilizing global temperature change until CO_2 concentrations in the atmosphere have been reduced. (Boucher et al., 2009) used a similarly simple climate model and carbon cycle model to show that when SRM was used to delay mitigation in this way, it must be sustained for several centuries in order to make up for an emissions overshoot of just a few decades. Matthews and Caldeira (2007) was the first paper to demonstrate the rapid climate change effects that would occur if SRM were started and then stopped abruptly (without making the requisite reductions in atmospheric greenhouse gases).

Results from GCMs

General circulation models (GCMs) are three-dimensional models of the atmosphere and/or ocean that simulate climate systems based on equations of balance of mass, momentum and energy; radiative transfer; and representations of cloud, land surface and ice processes. Atmosphere-ocean GCMs (AOGCMs), which include dynamic ocean and atmosphere systems that interact with each other at regular time intervals, are the most detailed and physically-based models of the climate system available today. Even AOGCMs, however, have significant limitations, in particular, in reproducing regional precipitation. (IPCC, 2007a) GCM simulations of SRM activities conducted in the past decade have

some common elements, such as general global scale cooling effects that over compensate for anthropogenic warming in the tropics and under-compensate in the poles and weakening of the hydrological cycle at the mean-global level. (Bala et al., 2008; Jones et al., 2010; Lunt et al., 2008; Robock et al., 2008) Several other findings, such as a weakening of the El Nino-Southern Oscillation (ENSO) (Lunt et al., 2008), a weakening of the African and South Asian monsoons (Robock et al., 2008) or effects of the thermohaline circulation (Caldeira and Wood, 2008; Lunt et al., 2008) have not been consistently observed in all models/experiments.

Global-scale Responses

A number of studies using AOGCMs confirm that SRM, whether implemented by injecting SO_2 into the stratosphere (Jones et al., 2010; Robock et al., 2008), or tuning down the sun (Bala et al., 2008; Caldeira and Wood, 2008; Lunt et al., 2008), can counteract rising global temperatures from GHGs. Note, Bala et al. (2008) and Caldeira and Wood (2008) both use AGCMs with slab oceans. All these studies also consistently show a weakening of the global hydrological cycle. Overall SRM tends to overcool the tropics relative to the poles (i.e., weaken meridional heat transport).

Regional Responses

One of the earliest studies examining SRM using a general circulation model, but with a slab ocean, Govindasamy and Caldeira (2000), showed SRM countering the temperature anomalies associated with anthropogenic climate change in an almost globally uniform way. Subsequent works, making use of more sophisticated models, have revealed more complex geographic disparities.(Jones et al., 2010; Lunt et al., 2008; Ricke et al., 2010; Robock et al., 2008) The specific regional disparities in temperature and precipitation response are not consistent between models. While there is regional diversity in the re-

sponse, a high-CO_2 world with SRM looks more like a low CO_2 world than a high CO_2 world without SRM. (Caldeira and Wood, 2008; Lunt et al., 2008)

Sea Ice & Applying SRM over the Arctic

One problem that SRM could be deployed to address is diminishing summer sea ice extents. Because of the strong positive feedbacks associated with diminishing sea ice cover, mitigating this effect of global warming could have a disproportionate impact on overall global climate changes. While most modeling studies that have looked at the effect of SRM on sea ice have found that implementing it globally does increase or stabilize Arctic sea ice cover, though this is not true in the experiment in Lunt et al. (2008). Several studies have explored the effects of deploying SRM over the Arctic or poles only. (Caldeira and Wood, 2008; Robock et al., 2008) They have found that limiting SRM forcing to over the Arctic does not result in cooling over the Arctic only, but that such schemes could be more effective than globally uniform ones in increasing Arctic sea ice cover, and may result in a modest decrease in meridional water vapor transport. (Caldeira and Wood, 2008)

Other Effects

Evidence of weaker African and Asian monsoons after major volcanic eruptions has raised concerns about whether SRM could have a similar effect. Robock et al. (2008) found evidence in its GISS-ModelE simulations of SRM, but other studies, such as those conducted with the Hadley Centre models, HadCM3L or HadGEM have not. (Jones et al., 2010; Ricke et al., 2010) Lunt et al. (2008) observed a weaker ENSO in their simulations of SRM, an effect that has not been explored in depth in other studies yet, but the lower resolution of HadCM3L over the ocean may make indicators about ENSO behavior simulated in that model less reliable.

Research Needs

While simulations of SRM have been done using a number of different coupled climate models, the experimental designs of the studies above are not consistent and therefore results from different models are difficult to compare directly. As a result, scientists from several major climate modeling groups are coordinating a series of SRM simulations, called the Geoengineering Model Intercomparison Project (GeoMIP) (Kravitz et al., 2011). GeoMIP will simulate four standard SRM scenarios using as a part of the Coupled Model Intercomparison Project (CMIP5) suite of experiments for the next IPCC Assessment Report. The work presented in Chapter 4 of this thesis, which explores parametric as opposed to structural uncertainty associated with modeling SRM, is complementary to this initiative.

1.5 Global Governance of SRM

SRM presents two global governance problems that are related, but different in character and scale. First, SRM is an emerging technology with a growing body of research on its potential implementation and effects, but there are still many associated risks and uncertainties. SRM research is, at this point, largely uncoordinated and unregulated. Norms and institutions are needed to ensure that SRM testing is safe, transparent and in support of an equitable approach to solving the global climate change problem. Second, SRM is a potential game changer in terms of climate change policy negotiations.

There is every indication from research thus far that this technology could stabilize global temperatures for a cost that is much less than mitigation (McClellan et al., 2010), making it an attractive but highly imperfect substitute or supplement for emissions reductions. In addition, its effectiveness in reducing the impacts of rising greenhouse gases varies by region and is uncertain. These unique attributes of SRM could shift various ac-

tors' preferences and strategic approaches toward climate policies. Despite the fact that there are some potentially grave impacts associated with SRM schemes, such activities are inexpensive enough that a number of nations could pursue such activities unilaterally. There is currently no international governance structure under which these activities can be restricted or monitored. Chapter 5 explores some of these issues in depth.

1.6 Overview of this Thesis

The remainder of this document is structured as follows:

- Chapter 2 covers the experimental set-up for the two sets of simulations of stratospheric SRM deployment executed using climate*prediction*.net;
- Chapter 3 presents the results of the first experiment on regional response to different levels of SRM;
- Chapter 4 presents the results of the second experiment on SRM efficacy under model uncertainty;
- Chapter 5 analyzes about how SRM changes the international relations of climate change and governance issues related to SRM research and deployment; and
- Chapter 6 presents some final thoughts and suggestions for future research.

Chapter 2

Methods

A multitude of models exist for examining the global climate today and into the future. Choosing a climate model for a given research initiative involves trade-offs between the strengths and limitations of these different types of models. The tradeoff when selecting a model is between complexity and computing time. The more complex a model is—for example, the higher its spatial resolution or the more processes occurring in each cell in each time step—the more computing power a single simulation requires. Lower complexity models allow the examination of longer time periods and/or larger parameter spaces for the same amount of computing power. Among the most complex climate models used for evaluating problems related to global warming issues today are coupled atmosphere-ocean general circulation models (AOGCMs). These models include dynamic ocean and atmosphere systems that interact with each other at regular time intervals.

2.1 HadCM3

The Hadley Centre Coupled Model, version 3 (HadCM3) is a coupled atmosphere-ocean general circulation model (AOGCM) that was developed by the UK MetOffice. It was one

of the models used for making global climate projections in the IPCC's third and fourth assessment reports. HadCM3 has 19 vertical levels in the atmosphere and 20 vertical layers in the ocean (with higher vertical resolution near the surface) and it progresses in half hour time steps. (Gordon et al., 2000) For the simulations presented in this document were run using HadCM3L, a version of HadCM3 reduced with resolution of 2.5° in latitude by 3.75° in longitude over the whole globe, as opposed to, in HadCM3, 2.5° in latitude by 3.75° over land, but 1° by 1° over the ocean. The exact specifications of the climate*prediction*.net(cpdn) version of HadCM3L used are identical to that used for the cpdn BBC Climate change experiment. (Frame et al., 2009)

2.2 climate*prediction*.net

The climate*prediction*.net project, based out of the Department of Physics at the University of Oxford, uses personal PCs around the world to run large ensembles of model variants ("perturbed physics ensembles") and assemble a distribution of climate responses to emissions scenarios, using different combinations of input parameters. Individual climate simulation work units are distributed to cpdn participants around the world using the Berkeley Open Infrastructure for Network Computing (BOINC) client, a grid computing tool originally developed for running the SETI@home project.(Stainforth et al., 2002) By pooling the resources of tens of thousands of computers around the world, rather than relying on a limited number of super-computers, cpdn has generated more results than any other climate-modeling project. (Allen, 1999; Frame et al., 2009). The cpdn version of HadCM3L allows perturbation of approximately 70 model parameters and forcing files to account for physical uncertainties associated with the standard configuration of the model. See Appendix A for a descriptions of the 22 parameters that were perturbed as a part of either of the experiments presented in this document.

2.3 Simulations

The experiments presented in Chapters 3 and 4 were designed to explore two different types of uncertainty in the climate response to stratospheric SRM forcings. Experiment 1 aims to see how global and regional responses to SRM vary with the amount of SRM applied. It used a single, standard configuration of the model to test 135 different SRM scenarios, revealing a variety of regional sensitivities to incremental changes in these forcings. Experiment 2 was designed to explore how climate response to SRM depends on the parametric configuration of the model. It examines the response of 43 different model configurations (parameter sets) to a much smaller set of SRM scenarios. There are a number of similarities in the design and implementation of the two experiments, however, and the common elements are described below.

Basic Set-up

All of the simulations conducted for the two experiments were transient, i.e., atmospheric composition changed over their length. Future anthropogenic greenhouse gas and tropospheric aerosol emissions were modeled using the IPCC Special Report on Emissions Scenarios (SRES) scenario A1B, a standard baseline emissions scenario for climate modelers which is nominally based on:

> ...a future world of very rapid economic growth, global population that peaks in mid-century and declines thereafter, and the rapid introduction of new and more efficient technologies. Major underlying themes are convergence among regions, capacity building and increased cultural and social interactions, with a substantial reduction in regional differences in per capita income...[and a] technological ... balance across [fossil and non-fossil fuel] sources. Nakiçenoviç and Swart, 2000

While SRES A1B is a helpful emissions scenario to use in these simulations because of its ubiquity in other climate modeling studies, the SRES projections are not necessarily very realistic (Morgan and Keith, 2008). After the first experiment was completed, a new set of emissions scenarios was published for use in the upcoming IPCC Fifth Assessment Report. The "Representative Concentration Pathways" (RCP scenarios) are intended to reflect improved understanding of the climate system and the likely pathways for human and economic systems that will influence it. (Moss et al., 2010) Scenario RCP 4.5 is approximately equivalent to SRES A1B (in its assumptions about economic growth, population and energy systems evolution, as well as its first-order forcing profile), however, Experiment 2 uses A1B for consistency and for reasons of expediency.

The simulations run from 2000-2080 with SRM forcings applied starting in 2005. Typical cpdn simulations run from 1920-2080, initiated with 150 year spin-ups. (Frame et al., 2009) Since the two SRM experiments are identical to other cpdn simulations, except for the SRM forcings starting in 2005, it was preferable to restart old simulations in 2000 rather than reinitiate from 1920. In order to initiate these 80-year simulations, we used restart files for the year 2000 from old simulations that had used identical parameters (and forcings through year 2000). For Experiment 1, we had only ocean restart files available and used the 1920 atmosphere start file. As such, we observed a blip at the beginning of the simulations as the atmosphere equilibrated with the ocean. For Experiment 2, we used only model configurations that had both ocean and atmosphere restart files available for the year 2000.

SRM Forcings

Stratospheric SRM activities were mimicked in the model by modifying the natural volcanic forcing inputs, which are implemented as zonally-uniform variations in strato-

spheric aerosol optical depth at 0.55 microns. The stratosphere in the model corresponds to the top 5 vertical layers of the atmosphere and the aerosol mass is distributed proportional to the air mass at each level. (Yamazaki, 2008) In both experiments, both SRM and no-SRM scenarios used globally uniform stratospheric forcings. Forcings were converted to stratospheric optical depth values using UK Met Office diagnostics. This approach to modeling SRM is unique and could be considered an intermediate approach between other recent SRM GCM modeling experiments which have either injected SO_2 into a dynamic stratosphere (Jones et al., 2010; Robock et al., 2008) or just directly tuned down the solar constant (Caldeira and Wood, 2008; Lunt et al., 2008).

HadCM3 was found to produce a forcing of approximately -2.5 W/m^2 for every 0.1 units of stratospheric optical depth. (Ingram, 2008) I assume that radiative forcing is linear with optical depth in the design of SRM scenarios, an imperfect assumption (Hansen et al., 2005), but adequate for the purposes of this experiment in which approximate forcing compensation was acceptable. The SRES A1B simulations with no SRM (our control) use a volcanic forcing file with a constant optical depth of 0.01, a value approximately equivalent to mean volcanic activity of the recent past (Sato et al., 1993).

Initial Condition Subensembles

In both experiments, all model versions were tested using ten-member initial condition subensembles. These ensembles, made up of simulations identical but for small perturbations to initial conditions, are important for reducing signal-to-noise issues associated with complex climate models. All AOGCMs, including HadCM3 have "internal variability," or noise associated with the internal dynamics of the model as opposed to the actual physical processes the model is simulating (Collins et al., 2001). Because of internal variability, otherwise identical simulations (with identical parameters and forcings) that

have small differences in their initial conditions will have very different states in any given time-step. By averaging the results of an initial conditions ensemble, it is possible to eliminate some of the noise of internal variability that interferes with the signals associated with the models response to natural and anthropogenic forcing inputs. The experiments use two mechanisms for generating initial conditions:

- The dtheta parameter: a temperature perturbation to one grid cell; and
- Simulation restart files: as described in the section above, simulations for this experiment were restarted in model-year 2000 using ocean and atmosphere conditions saved from previous simulations. Using different restart files for the same model configuration is another way to perturb our year 2000 initial conditions.

For Experiment 1, only dtheta perturbations were used to generate initial condition ensembles. In Experiment 2, when more than one simulation for a given model version was available, semi-random combinations of restarts and dtheta perturbations were used to generate the initial condition ensembles. (Some additional details are available in Chapter 4.)

Regionalization

To systematically analyze the regional implications of the results temperature and precipitation anomalies are analyzed grouped into in 23 regions output by climate*prediction*.net. The regions are shown in Figure 2.1 and described in Table 2.1. These "Giorgi" regions, were designed to represent climatically and physiographically similar land areas and are large enough to produce climate predictions that are more statistically robust than those obtained from grid-cell level output (Giorgi and Francisco, 2000). Population and economic

output data for the year 2000, also shown in Table 2.1, was mapped to the regions using the Nordhaus (2006) G-Econ dataset.

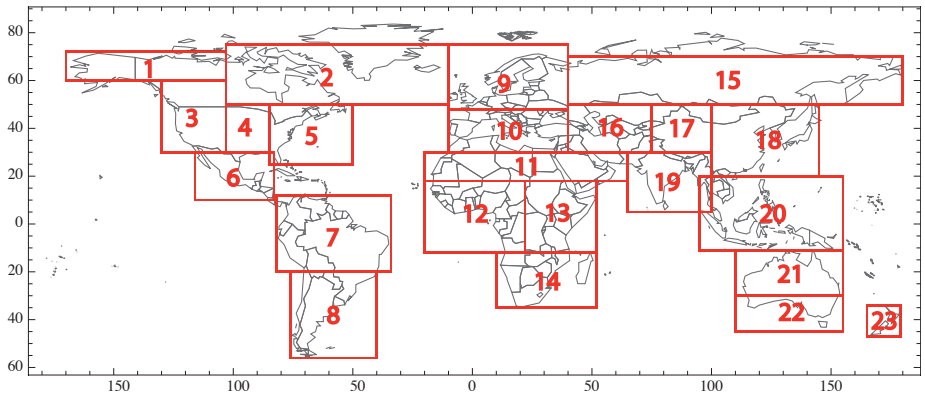

Figure 2.1: The 23 Regions used in the analyses.

2.4 Analytical Tools

The following tools were used for simulation data management, processing and presentation:

- The CPDN download client: the Mac OSX CLI binary, available at https://results.cpdn.org/repository/help, was used to download zip files of simulation output from the CPDN servers.

- Python: this programming language was used to write scripts for data management (unzipping and moving output files) and data processing (opening netCDF files, extracting relevant data and writing into txt or csv output files). The netCDF4 module (Dataset class) was used for reading cpdn netCDF files. Some simple statistical processing (e.g., averaging initial condition ensembles, calculating standard deviations of baseline datasets) was conducted using NumPy functions.

Table 2.1: Regional Data

No.	Region	Long Name	Economic Output (billions US$)	Population (millions)
1	ala	Alaska, NW Canada	40	0.58
2	cgi	E Canada etc.	22	0.90
3	wna	Western N America	2592	74.67
4	cna	Central N America	2659	82.75
5	ena	Eastern N America	5149	148.36
6	cam	Central America	1319	134.29
7	amz	Amazonia	1211	201.83
8	ssa	Southern S America	1119	150.52
9	neu	Northern Europe	7175	378.39
10	seu	S Europe, N Africa	5329	449.52
11	sah	Sahara	730	70.67
12	waf	Western Africa	304	299.42
13	eaf	Eastern Africa	215	267.37
14	saf	Southern Africa	356	121.91
15	nas	Northern Asia	697	84.28
16	cas	Central Asia	744	288.50
17	tib	Tibetan Plateau	153	63.79
18	eas	Eastern Asia	7105	1485.25
19	sas	Southern Asia	1784	1271.44
20	sea	Southeast Asia	1420	470.08
21	nau	Northern Australia	109	4.36
22	sau	Southern Australia	401	14.80
23	nz	New Zealand	71	3.86

- Mathematica: this software package was used for most of the data analyses and figure generation. The CountryData function was used for mapping. The LinearModelFit function was used for regression analyses.

- R/ RStudio: the statistical programming language, R, as implemented through the RStudio environment, was used for generating regression trees using the rpart (recursive programming) package.

- Adobe Illustrator: this graphics editing software was used to create graphics and to improve the readability and aesthetics of figures generated in Mathematica or RStudio.

Chapter 3

Results: Experiment 1

This chapter presents an analysis of regional differences in climate modified by SRM using a large-ensemble modeling experiment that examines the impacts of 54 global temperature stabilization scenarios. Our results confirm that a world with SRM would generally have less extreme temperature and precipitation anomalies than one with unmitigated greenhouse gas emissions and no SRM, but illustrate the physical unfeasibility of simultaneously stabilising global precipitation and temperature as long as greenhouse gases continue to rise.

Over time, simulated temperature and precipitation in large regions such as China and India vary significantly with different SRM trajectories and diverge from historic baselines in different ways. Hence the use of SRM to stabilize climate in all regions simultaneously may not be possible. Regional diversity in the response to different levels of SRM could complicate what is already a very challenging problem of global governance, and could make consensus about the "optimal" level of geo-engineering difficult, if not impossible, to achieve. [1]

[1] This chapter is based on an article written with M. Granger Morgan and Myles Allen that was published in the August 2010 issue of *Nature Geoscience*. (Ricke et al., 2010) Myles and I designed the experiment, I performed all the data analyses and drafted the article and Granger and Myles contributed to refining it.

3.1 SRM Scenario Design

135 SRM scenarios were formulated, designed to offset the net anthropogenic forcings under SRES A1B. The major anthropogenic climate forcings included in the cpdn-version of HadCM3 are from long-lived greenhouse gases, tropospheric ozone and sulfur aerosols (direct and first indirect effect). In the first step in designing counterbalancing geoengineering scenarios, the transient magnitudes of these forcings were quantified and summed. The forcing scenarios were based on the greenhouse gas concentrations and tropospheric ozone and sulfur aerosol burdens derived for the SRES A1B emissions scenario— based on a future with high economic growth, low population growth and rapid technological change. (See Chapter 2 for details.) The baseline values over time of the three forcing components are shown in Figure 3.1, along with the uncertainty ranges applied according to the method below.

Greenhouse gas concentrations for SRES A1B were taken from the IPCC Third Assessment Report (TAR) and relative forcings were modeled using the approximation equations in Table 6.2 of the TAR WG1 Report. (Ramaswamy et al., 2001) Gases included are carbon dioxide, methane, nitrous oxide, CFC-11 and CFC-12. The aggregate greenhouse gas forcing grow from 2.4 W/m2 in 2000 to 5.7 W/m2 in 2080. The uncertainty associated with the forcings from long-lived greenhouse gases are approximately ±10% from IPCC Fourth Assessment Report (AR4).(Forster et al., 2007)

Decadal values of tropospheric ozone burdens were obtained from Horowitz (2006). The burdens used to scale the ozone forcing relative to a year-2005 reference forcing estimate and concentration. The estimated radiative forcing effect of tropospheric ozone in

The linearity analysis at the end of the chapter was presented in a poster at the 2009 AGU Fall Meeting and further refined and applied as part of the Residual Climate Response (RCR) model analysis. The RCR model application is presented and summarized in the final section and is adapted from a paper written with Juan Moreno-Cruz and David Keith that is in press at *Climatic Change*. (Moreno-Cruz et al., 2011) Juan and I performed the data analysis for this paper together.

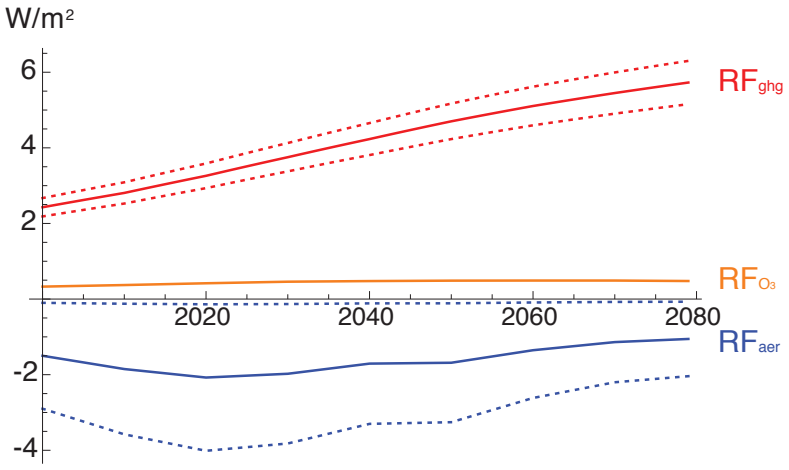

Figure 3.1: Anthropogenic Forcing Values Over Time. Greenhouse gas forcing estimates in red, sulfur aerosols in blue and tropospheric ozone in orange.

2005 was 0.35 (-0.1, +0.3) W/m2. (Forster et al., 2007) Due to both the small relative value of the tropospheric ozone forcing and the limited number of forcing schemes that could be tested in this experiment, ozone uncertainties were not specifically included in the design of the geoengineering forcing files. Rather, ozone forcing values were combined with the forcing estimates for long-lived greenhouse gases and varied and the combined spread of forcing values were tuned according to the methodology discussed in the next section.

To approximate the radiative forcing effects of sulfur aerosols, a sulfur burden index, relative to the year 2000 was developed using the dataset from Boucher et al. (2002). Direct and indirect sulfur aerosol effect forcings were estimated and applied over time using linear scaling. The baseline value used for the direct effect forcing is -0.4 W/m2, as reported in AR4, with 90% confidence intervals of ±0.2 W/m2. A baseline value for the indirect effect was established using the modeling studies in Table 2.7 of the AR4 WG1 report in which the model used estimated the effect of sulfur aerosols only. In total, five results

met this criterion and the mean radiative forcing value was -1.1 W/m2 with a standard deviation of 0.5 W/m2.

The spread of radiative forcing values associated with greenhouse gases and sulfur aerosols was tuned to fit the approximate spreads presented in Chapter 2 of the AR4 WG1 report on radiative forcing estimates, Figure 3.2.(Forster et al., 2007) Because the goal of this experiment was to identify the effects associated with optimal geoengineering schemes, the scenarios span the entire range of potential forcing values, even those that are highly unlikely. The first iteration of forcing files was produced by summing incremental variations in the greenhouse gas and aerosol forcing values and the resulting ensemble of forcing files was expanded and tuned to fit with these summary data.

In total, 15 sulfur aerosol forcing timelines were crossed with 9 greenhouse gas timelines for a total of 135 possible geoengineering forcing scenarios. Forcings were converted to stratospheric optical depth values using UK MetOffice diagnostics (Ingram, 2008). HadCM3 was found to produce a forcing of approximately -2.5 W/m2 for every 0.1 units of stratospheric optical depth. The 135 scenarios are pictured in Figure 3.3.

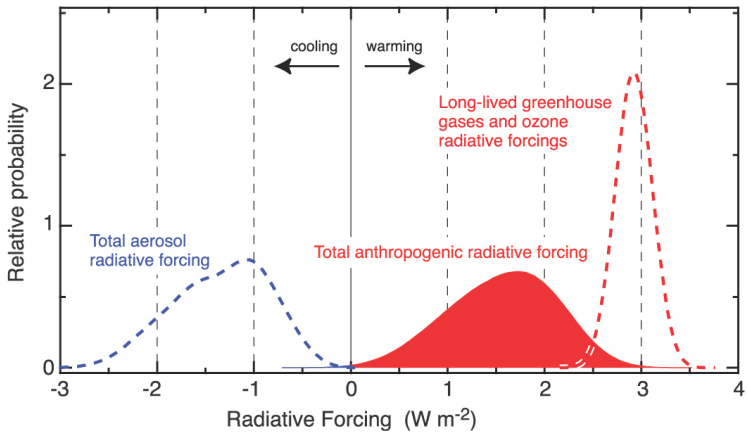

Figure 3.2: Figure 2.20B from Chapter 2 of the 2007 IPCC Working Group 1 Assessment Report: Probability distribution functions (PDFs) from combining anthropogenic radiative forcings. (Forster et al., 2007)

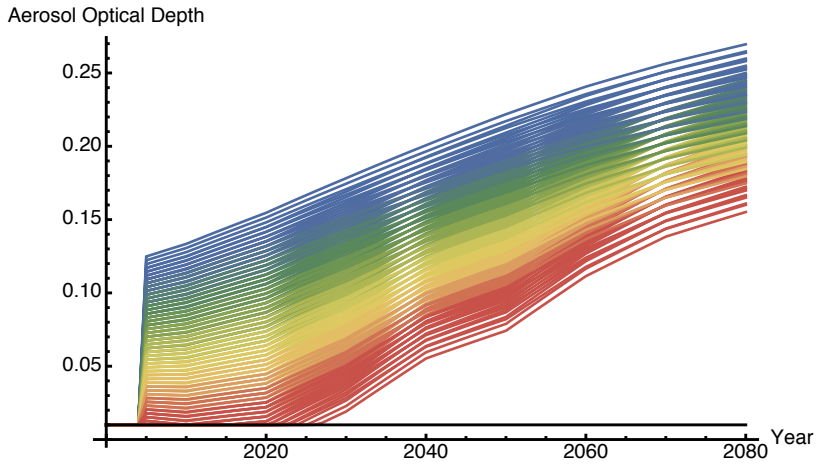

Figure 3.3: 135 SRM Forcing scenarios, implemented as set values of stratospheric aerosol optical depth.

3.2 Global Response to SRM Scenarios

Of the 135 SRM scenarios considered, a least-squares fit analysis was used to select 54 for which the global surface air temperature trend was less than $\pm 0.006^{\circ}$C per year (most SRM scenarios produced global temperature trends that rose slightly over the length of the simulation) (Figure 3.4a). All 54 of the SRM scenarios produced stabilized five-year average global-mean surface air temperatures, at levels between approximately 14.6 to 15.7°C (Figure 3.4b) – roughly, plus or minus half a degree from the temperature at the time SRM activities are initiated– depending on the level of forcing applied, while the control scenario (shown in black) resulted in an increase in global-mean SAT of approximately 2.5°C over the course of the 80-year simulations. For the control scenarios, seasonal temperature maps of the anomaly in surface air temperature between the 1990s (the last common decade of data for both sets of simulations) and the 2070s show warming everywhere, but especially at the poles during local winter. As presented in the maps in Figure 3.5, these effects are largely neutralized in the runs with SRM, although there is greater cooling in the tropics than elsewhere.

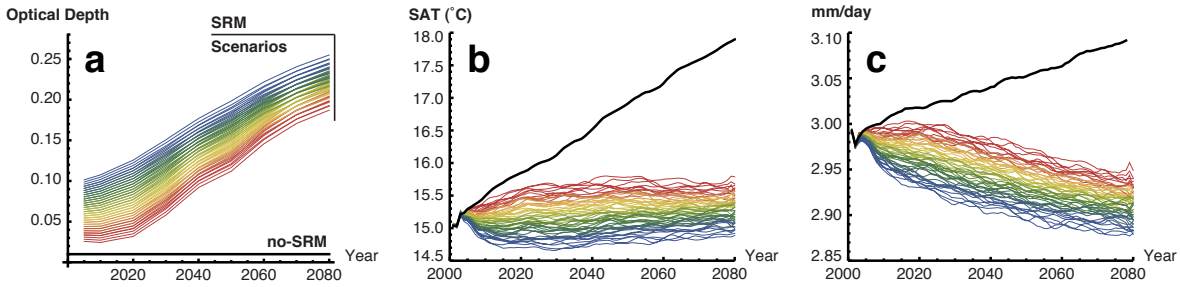

Figure 3.4: Time series of global optical depth, temperature and precipitation of the 54 scenarios examined. (**a**) Solar radiation management (SRM) and no-SRM scenarios, (**b**) five-year average global mean near-surface (1.5 m) air temperature, and (**c**) five-year average global mean precipitation rate, all displayed over the length of the 80 model-year simulations.

As theoretical frameworks (Allen and Ingram, 2002) and previous modeling results (Robock et al., 2008) have predicted, we find a global net increase in precipitation under the control (no-SRM) scenario and net decreases under the scenarios with SRM (Figure 3.4c). Because of the component of the hydrological impact of long-wave forcing that is independent of temperature, SRM with stratospheric aerosols cannot simultaneously compensate for the impacts of rising greenhouse gases on both temperatures and the hydrological cycle.(Bala et al., 2008) While it might be possible in principle to "fine tune" the hydrological response by injecting aerosols with different optical properties at different latitudes or altitudes, no proposal yet exists for how this might be implemented in practice, and some variability in response remains inevitable. Hence, as Figure 3.6 illustrates, uniform SRM cannot compensate exactly for rising greenhouse gas concentrations at the global level. The maps in Figure 3.7 show seasonal precipitation anomalies for the no-SRM and one SRM scenario. The geographical distribution of precipitation effects varies widely under both sets of simulations, with globally increased albedo sometimes mitigating the precipitation anomalies exhibited under the standard global warming scenario, but oc-

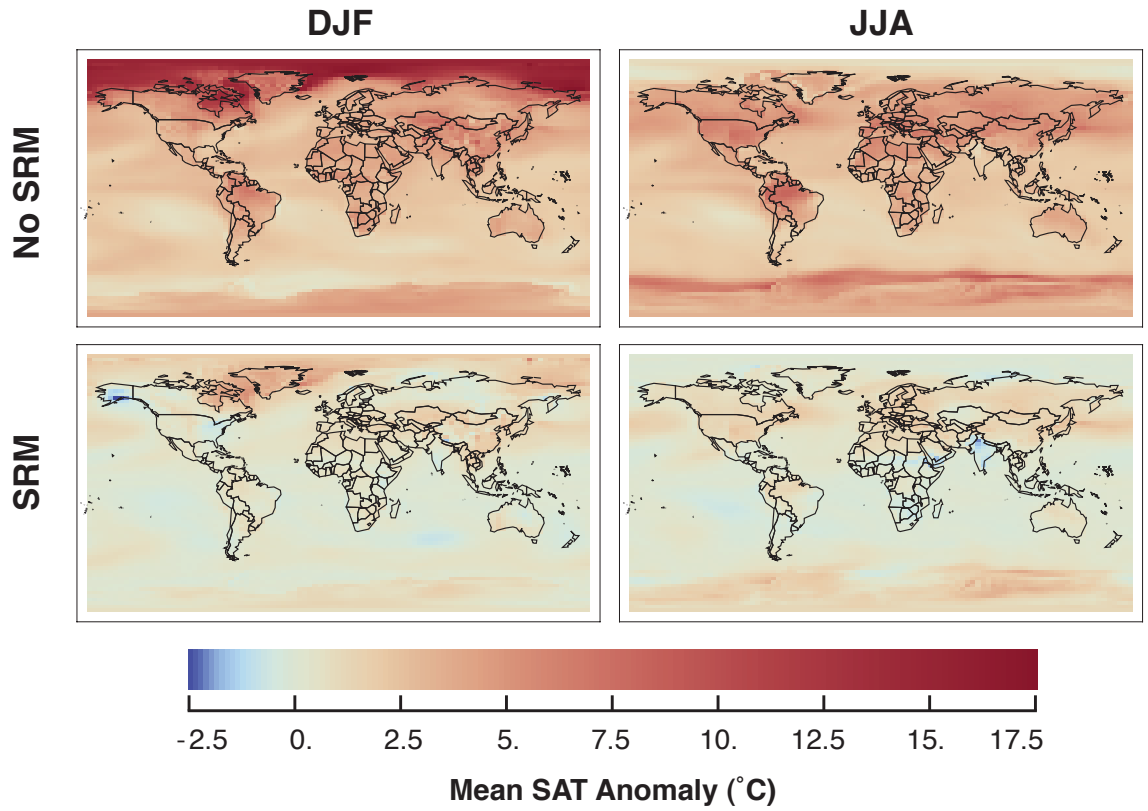

Figure 3.5: Near-Surface (1.5 m) Air Temperature Anomaly (<2070s>-<1990s>) Maps for simulations with no SRM (n=30) and SRM that stabilized global mean SAT at approximately 15.2°C (n=33).

casionally exacerbating them. Previous studies have not examined how global patterns of these changes vary with different SRM scenarios.

Surface and subsurface runoff anomalies are generally mitigated with SRM. Maps of seasonal subsurface runoff are presented in Figure 3.8. This indicator is an alternative proxy (beyond precipitation) for the effect of both scenarios on regional hydrological cycles because it is a function of soil properties, temperature and precipitation, and therefore may better portray information about impacts. Note that in the subsurface runoff data, the sole region for which runoff anomalies under the no-SRM scenario are not mit-

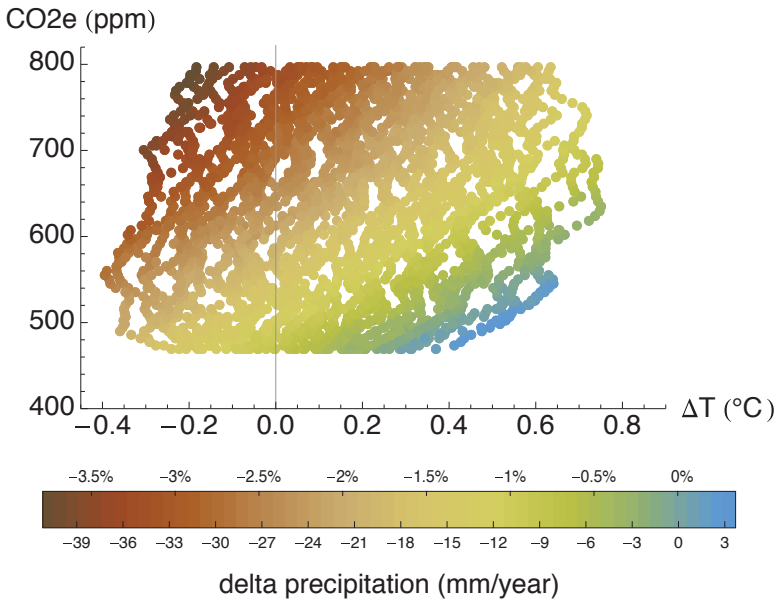

Figure 3.6: Relation between change in precipitation rate as a function of change in global mean near-surface air temperature (x-axis) and equivalent carbon dioxide concentration (y-axis). Temperature and precipitation data points are five-year averages over the initial-condition sub-ensembles for each SRM scenario.

igated under SRM scenario is Southeast Asia in JJA.

3.3 Regional Results

In order to analyze the regional implications of different levels of SRM we examined mean temperature and precipitation anomalies over land in 23 macro-regions (Giorgi and Francisco, 2000). A summary plot of regional temperature and precipitation responses to the different forcing scenarios early and late in the simulations is shown in Figures 3.12 and 3.13 (placed at the end of the chapter). While increased stratospheric albedo cools all regions considered compared with the A1B control, precipitation responses vary. In most regions, our simulations support the general assumption that the more SRM that is im-

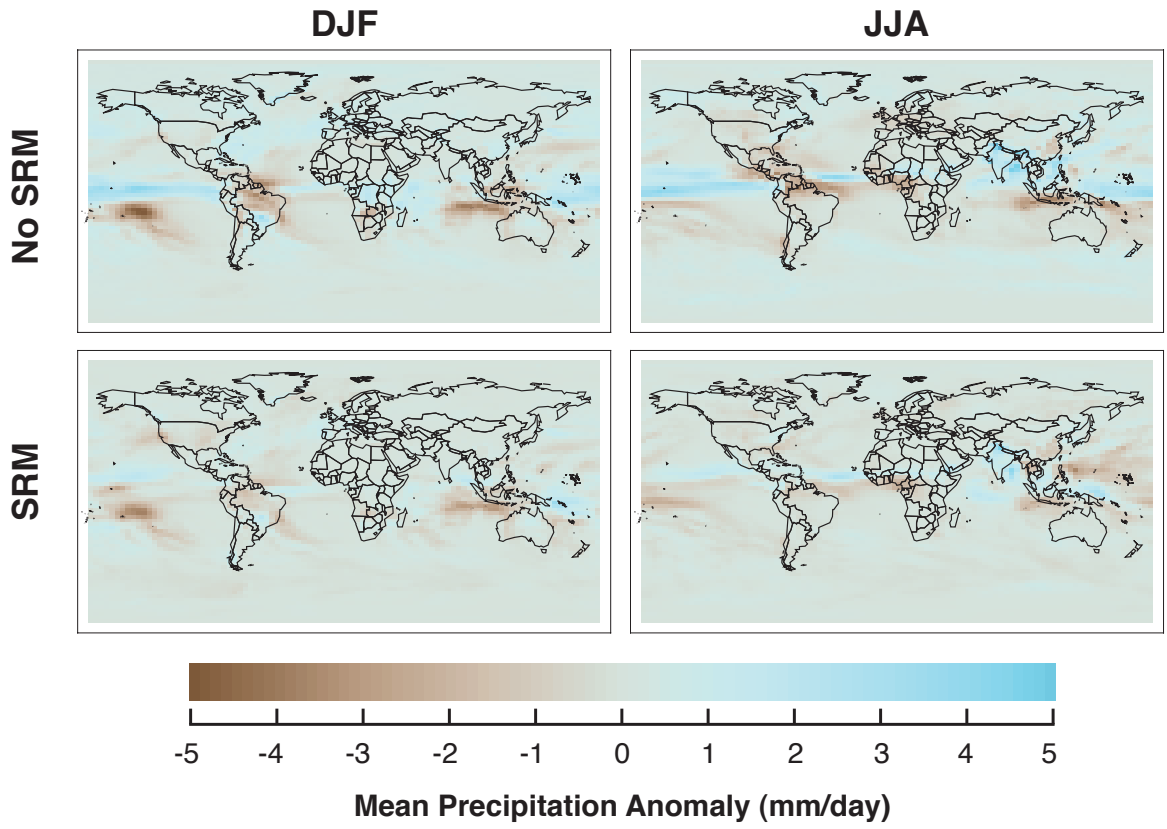

Figure 3.7: Precipitation Rate Anomaly (<2070s>-<1990s>) Maps for simulations with no SRM (n=30) and SRM that stabilized global mean SAT at approximately 15.2°C (n=33).

plemented, the greater the reduction in precipitation. However, there are some exceptions, such as Central America and the Amazon, and to a lesser extent, Southern Africa and the Mediterranean, although these regional details are likely to be sensitive to the model used. In most regions and seasons, there is a SRM scenario that produces precipitation rates closer to the baseline value than the control scenario, but again there are exceptions, such as Southeast Asia and Western North America in the summer. Precipitation in some regions, such as Canada and Northern Asia, is relatively sensitive in these simulations to the SRM scenario employed. Other regions, such as Australia and Eastern Africa are insensitive to the scenario.

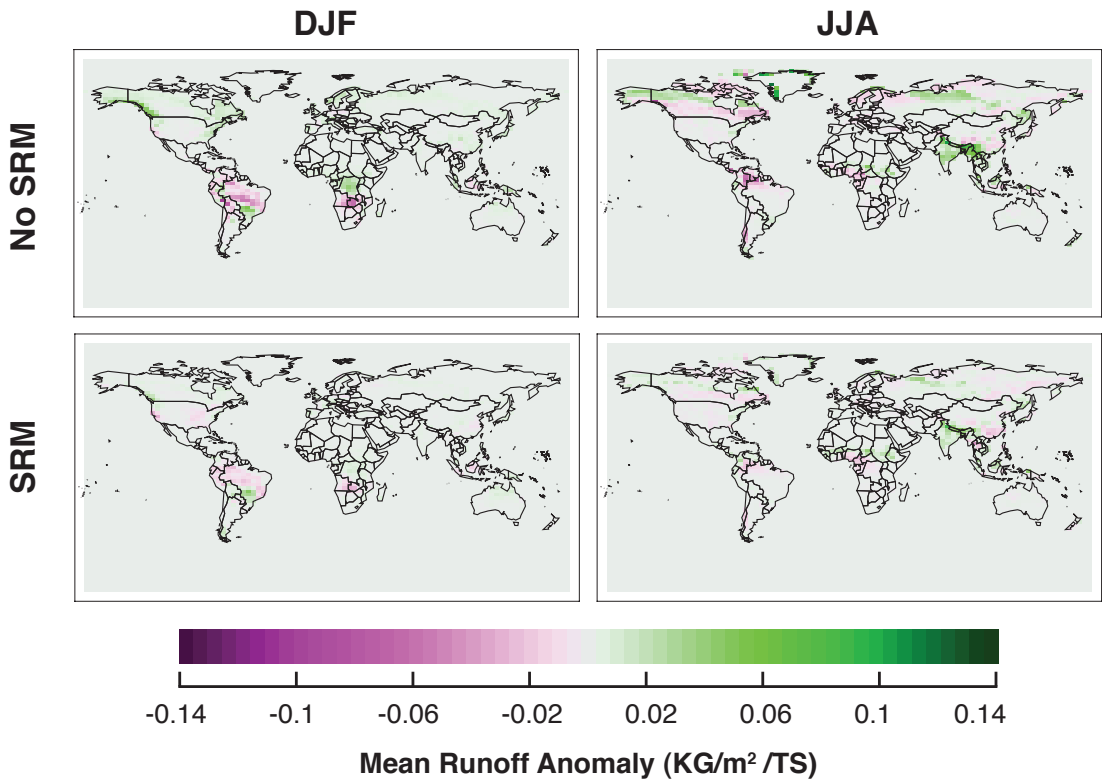

Figure 3.8: Subsurface Runoff Anomaly (<2070s>-<1990s>) Maps for simulations with no SRM (n=30) and SRM that stabilized global mean SAT at approximately 15.2°C (n=33).

The primary normalization scheme used to compare the anomalies observed in the 23 macro-regions analyzed is seasonal interannual variability from the baseline dataset (1990s transient standard physics simulations). This normalization scheme was selected as a recent common reference point that could provide a metric for how the general changes that happened over time in the SRM and no-SRM simulations compared to the (simulated) mean climate state and year-to-year variability of each region in the recent past. From an impacts relevance standpoint, the smaller the changes to the mean climate of a region are compared to interannual variability in the recent past, the easier such changes will be to adapt to. Other approaches to normalization can yield somewhat different results,

however the main conclusions are robust. For a further discussion of normalization, see Appendix B.

If the aim of SRM is to "restore late-20th-century climate", one way of defining this target would be to return a region's average temperature and precipitation to within one standard deviation of its baseline climate. Early in the simulations, a variety of SRM scenarios achieve this (see Figure 3.12 at the end of the chapter). However, by the end of the simulations, when SRM has compensated for increasing anthropogenic forcings for six to seven decades, there is often no scenario that can place a region back within one standard deviation of both its baseline temperature and precipitation (see Figure 3.13 at the end of the chapter). In other words, as the level of modification required to compensate for anthropogenic greenhouse gas forcing increases, the relative appeal of different levels of SRM depends on the region considered and the variable (temperature or precipitation) that is deemed most important.

This point is illustrated clearly in Figure 3.9, which shows the summer temperature and precipitation anomalies in the 2020s and 2070s for the regions containing India and Eastern China, normalized by the ensemble-mean interannual variability of their baseline (late-20th-century) climates. In the 2020s most scenarios return the climate of both Eastern China and India to within a standard deviation of baseline. By the 2070s, the scenarios that return regional climates to within the one-standard deviation circle are mutually exclusive. However, because of the large temperature anomalies in the no-SRM scenario by the 2070s, the net regional temperature-precipitation anomalies in these and other regions under the no-SRM scenario are much larger than those in any of the SRM scenarios.

Figure 3.10 shows the level of SRM that brings the regional climate back closest to its 1990 climate (i.e. to the centre of a circle like the one shown in Figure 3.9). That is, the amount of SRM that minimizes combined temperature and precipitation anomalies in

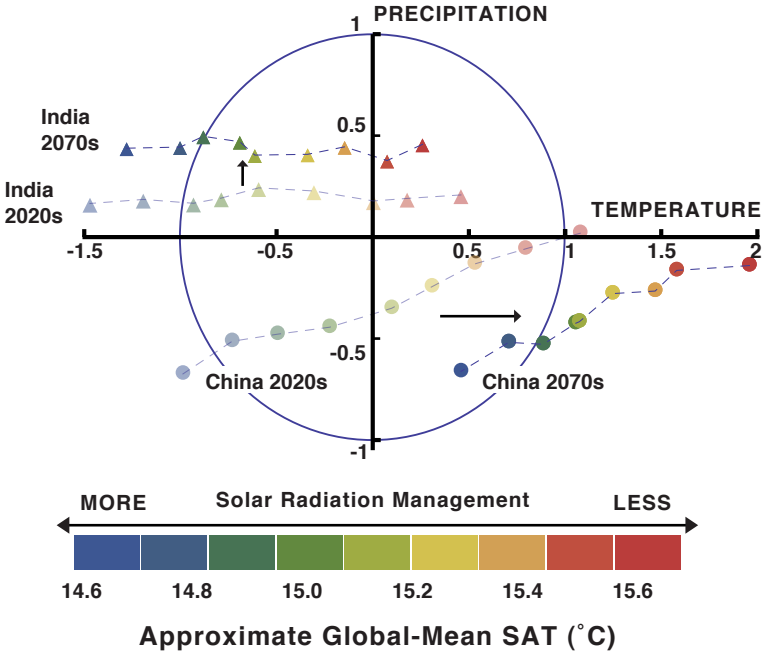

Figure 3.9: Modeled response to different levels of average global SRM over time in India and China. Interannual variability-normalized regional temperature and precipitation summer (JJA) anomalies (averages for the 2020s minus the 1990s and 2070s minus the 1990s) in units of baseline standard deviations for region including India (triangles) and region including Eastern China (circles). SRM-modified climates for these two regions migrate away from the baseline in disparate fashions.

units of regional baseline standard deviations. Again, we find that different levels of SRM will likely be desired by different regions.

3.4 Discussion

HadCM3 has been shown to reproduce the observed water vapour feedbacks associated with volcanic eruptions (the proxy process we use to mimic the SRM forcings) fairly well (Forster and Collins, 2004). However, it is well known that present GCMs are limited in their ability to estimate local and regional precipitation (Randall et al., 2007) and circu-

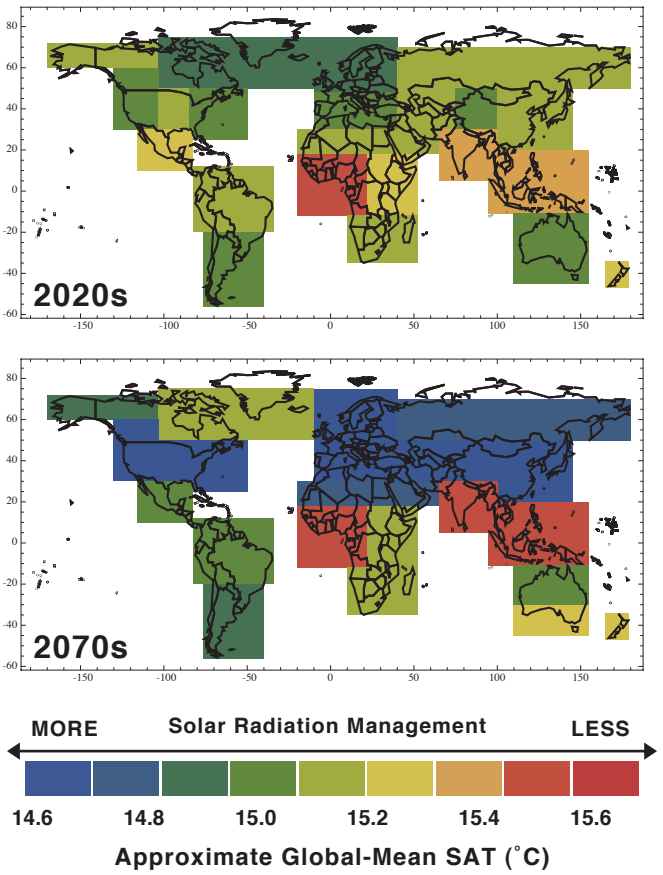

Figure 3.10: "Optimal" SRM scenarios for the summer for each region. (For regions that straddle the equator, the wetter season was selected.) The "optimal" scenario is defined as minimum combined temperature and precipitation anomalies in standard deviations from the baseline (1990s) climate, for (a) the 2020s and (b) the 2070s.

lation response in models of this class to volcanic forcing may be imperfect (Stenchikov et al., 2006). Some have proposed the application of non-uniform SRM forcings, such as stratospheric forcings concentrated over the poles, or even tropospheric forcings with both latitudinal and longitudinally varied distribution over the oceans, to control the regional impacts of global cooling by SRM (Latham et al., 2008). SRM implemented as such would likely produce different regional precipitation effects as well. However, while the specific patterns might be somewhat different under any actual implementation of SRM,

none of these interventions will have purely local impacts and it seems most unlikely that some form of the regional divergence that we observe would not appear.

Previous modeling exercises have demonstrated that unless net carbon dioxide emissions are reduced close to zero, with substantial reductions in emissions of shorter-lived greenhouse gases, long-wave forcing would continue to rise and SRM would be needed to compensate for global warming for centuries before it could be phased out (Boucher et al., 2009). However, as our simulations progress, regional geoengineered climates migrate away from the baseline origin. While all of the regions are closer to their 1990 conditions (especially in terms of temperature) under most of the scenarios, the most desirable level of SRM will very likely become more and more dependent on the region and variable considered the longer these activities are carried out as long-wave forcing continues to increase.

While the analysis of this chapter has examined mean temperature and precipitation anomalies over land in 23 macro-regions, it is important to note that these are not the only metrics that might be used in designing an SRM intervention. For example, in some regions sustaining annual or seasonal water resources (e.g. by protecting snow pack, or by assuring the continued operation of a monsoon system), or retaining summer sea ice, might be chosen as more important (and perhaps conflicting) objectives. Even within a region there may often not be an agreed metric. For example, indigenous peoples may want to preserve the summer arctic sea ice that is critical to their way of life, while maritime industries may prefer summers with an ice-free arctic ocean.

Even if the specific aims of SRM were agreed, predicting this kind of detail in the direct and indirect consequences of SRM is generally beyond the capabilities of the models currently used for SRM research. SRM at levels considered in this paper may result in a variety of unexpected and unintended consequences. For example, aerosols from the eruption of Mount Pinatubo produced more diffuse light through scattering and in the year

following the eruption a deciduous forest in Massachusetts experienced a 23% enhancement in photosynthesis because the plant canopies use diffuse light more efficiently than direct sunlight (Gu et al., 2003). Some stratospheric particles can provide reaction sites for the catalytic destruction of stratospheric ozone (Tilmes et al., 2009). If SRM were started and then stopped, while greenhouse gas concentrations continued to increase, the result would be unprecedented rates of rapid warming (Matthews and Caldeira, 2007) which have the potential to prove devastating to many terrestrial ecosystems. Of course, SRM alone would do nothing to stop the rise in of atmospheric concentration of carbon dioxide which will have differential impacts on species within terrestrial biota (Körner and Bazzaz, 1996) and lead to continued ocean acidification and the likely demise of many or all coral reefs. (Hoegh-Guldberg et al., 2007)

These results do not provide a definitive illustration of regional climate impacts associated with potential future SRM schemes, nor can we assume that minimizing the net normalized temperature and precipitation anomalies from the 1990s would be the objective of any given region should SRM be undertaken, although some such quantitative metric would be required to define the objectives of any SRM intervention. Rather, our results demonstrate that not only would "optimal" SRM activities imply different things for different regions, but that international negotiations over the amount of SRM could become inherently more difficult the longer such activities are used to compensate for rising greenhouse gas concentrations. While greenhouse gas emissions result from economic activity all over the world, the intentional modification of albedo could be undertaken by just one or a few parties. Consideration by diplomatic and other communities of how global governance of such activities might best be managed is in an even earlier stage of development (Victor et al., 2009). Results presented here suggest that as our understanding improves, serious issues of regionally diverse impacts and inter-regional-equity may further complicate what is already a very challenging problem in risk management and

governance.

3.5 Linearity of Regional Responses

From Figure 3.9, incremental changes in the amount of SRM appear to result in approximately linear temperature and precipitation changes, even at the regional level. A strong linear response of regional temperature and precipitation to changes in stratospheric aerosol optical depth (AOD – i.e., the amount of SRM) in the forcing range of interest is important because it means that in future experiments (i.e., Experiment 2 in the next chapter), far fewer SRM scenarios need be simulated in order to get a clear picture of regional sensitivities. With sufficiently large initial condition ensembles (to eliminate noise), the regional responses between high and low levels of SRM can be interpolated. A least-squares regression tested a linear approximation of temperature (or precipitation) given AOD:

$$\Delta T(od) = \alpha_1 + \beta_1 \cdot od$$

versus a function that includes a quadratic term:

$$\Delta T(od) = \alpha_2 + \beta_2 \cdot od + \gamma \cdot od^2$$

to determine whether it substantially improves the explanatory power of the model, assuming that if it does not, a low adjusted-R^2 ($\bar{R}^2$) is due to noise rather than changes in the indicator following a more complex functional form. Appendix C presents tables with the regression coefficients and $\bar{R}^2$ values associated with annual, Northern Hemisphere summer (JJA) and Northern Hemisphere winter (DJF) decadal mean temperature and precipitation anomalies for the 2020s and the 2070s. Given the high $\bar{R}^2$ values for the regional temperature-AOD models (always greater than 90% in the 2020s and greater

than 70% 2070s, for all regions and seasons), a linear approximation appears to be a good assumption.

Even with decadal means and 10-member initial condition ensembles, $\bar{R}^2$ values for precipitation data are quite low (and sometimes negative) for many of the regions, but the explanatory power of the model is rarely improved by adding a quadratic term. As such, the question is how much drastically reducing the number of SRM scenarios tested is likely to alter the results we get if more scenarios were empirically tested. I compared best-fit lines through two endpoints (one of the five highest SRM scenarios and one of the five lowest) and one randomly selected point between them in all 25 combinations to see how different the results were from the best-fit found using all models for each region and season. On average, the best-fit between precipitation and AOD that is interpolated using three points instead of 54 points predicts annual-mean regional precipitation anomalies that were about one hundredth of a standard deviation different or less in the SRM forcing range of interest. This implies that the three scenario approach taken for the analysis of Experiment 2's simulations in the next chapter will not produce significantly different results than if many more SRM scenarios had been tested empirically, as long as the extrapolation is made within the same range of forcings tested in this experiment.

3.6 Application Example: Residual Climate Response Model

The data from the experiment above was used to test a simple model to account for the potential effectiveness of solar radiation management (SRM) in compensating for anthropogenic climate change. (Moreno-Cruz et al., 2011) The Residual Climate Response (RCR) model, described below, measures how effective SRM is in compensating for CO_2 equiva-

lent (CO_2e)-driven climate change at a regional level. The model can be applied to evaluate any climate indicator in which regional responses are approximately linear in the forcing range of interest. Note that for this application, a stricter linearity criterion must be met than that in the section above – one in which the SRM forcing "range of interest" extends to the point of no SRM. (Appendix D presents the linearity tests specific to this analysis.)

In the RCR model, regional inequalities are represented by treating the climate changes induced by CO_2e and SRM as independent vectors in a space of relevant climate variables. The changes in a given climate indicator relative to a baseline and after SRM compensation, can be captured by a vector of residuals. The larger the magnitude of this vector, the lower the effectiveness of SRM. To account for differences in interregional preferences in a way that is impacts-relevant, we weighted the variability-normalized regional temperature and precipitation changes using welfare indicator data to represent three different social objectives: egalitarian, where each region is weighted by population (number of people), utilitarian, where each region is weighted by its economic output (US$ billion), and ecocentric, where each region is weighted in terms of Area (km^2). (Nordhaus, 2008)

Figure 3.11 shows a two-dimensional example of the RCR model. Assume the climate variable of interest is temperature and consider changes in two regions, Region A and Region B with populations a and b, respectively. For each region we calculate the population weighted temperature changes due to an increase in CO_2e, denoted by $T^A_{CO_2}$ and $T^B_{CO_2}$. Now assume that SRM is designed and implemented to minimize deviations from some reference point, e.g. preindustrial levels or 1990s levels. The reference point is represented as the origin in Figure 3.11. The temperatures in regions A and B after implementation of SRM are given by T^A_{RES} and T^B_{RES}. Because $\varphi > 0$, the optimal level of SRM leaves Region A with a positive change in temperature and Region B with a negative change in temperature. Notice that this is globally optimal. That is, the level that minimizes squared deviations in both regions, is not the optimal level for each region. If Region B chooses the

level of SRM that minimizes its own damages, temperature changes in Region B would be zero and temperature damages in Region A would be higher than at T^A_{RES}. The opposite is also true if Region A chooses the level of SRM to minimize its own damages. In this case, the temperature in Region B would be higher than T^B_{RES}. This shows the implementation of an optimal level of SRM is very difficult because different regions want different levels of SRM. Calculating the globally optimal level of SRM, however, while theoretical, serves as a benchmark towards which we can compare other policies.

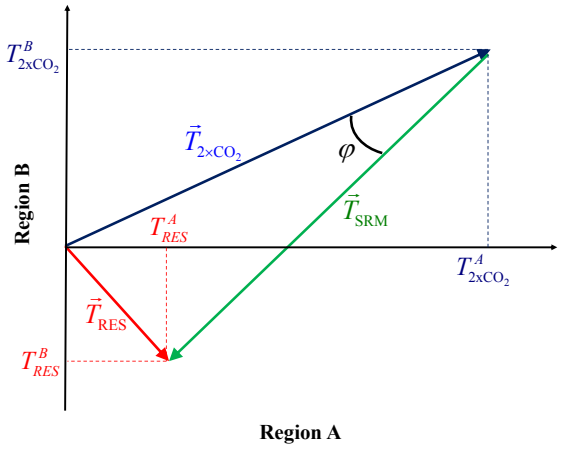

Figure 3.11: Residual Climate Response Model. The horizontal axis shows changes in temperature for Region A. The vertical axis shows temperature change in Region B. The blue vector represents the change in temperature for the two regions at the time of doubling CO_2e. The green vector represents the level of SRM that minimizes the sum of square temperature change in the two regions. The red vector shows the optimal level of SRM that minimizes the deviation from the baseline for the two regions simultaneously. The angle φ measures the effectiveness with which SRM compensates for CO_2e-driven temperature change. Under the common assumption that impacts are quadratic in temperature deviations, there is an equivalence between the length of the residual vector and the total damages after the implementation of SRM. This same logic applies for different climate variable (e.g. precipitation) and more than two regions.

We use the vector of residuals to obtain a more intuitive measure of compensation. In particular, we assume regional damages are an increasing function of changes in a given

climate variable (Nordhaus, 2008; on Stabilization Targets for Atmospheric Greenhouse Gas Concentrations) [Nordhaus 2008, NAS 2010]. For simplicity, we assume damages are quadratic and we estimate the percentage of climate change-related regional damages that can be compensated for using SRM. That is, damages D from a change in the climate indicator Y, can be approximated as follows:

$$D_{CO_2e} \propto \|\Delta \mathbf{Y}_{CO_2e}\|^2 \qquad \text{and}$$

$$\mathrm{D}_{RES} \propto \|\Delta \mathbf{Y}_{\mathbf{RES}}\|^2$$

Using this assumption, we can say that the percentage of damages from regional changes in the variable Y compensated for with SRM is given by the following equation; where the right hand side follows from the definition of $\mathbf{Y}_{RES}$:

$$\left(1 - \frac{\|\Delta \mathbf{Y}_{RES}\|^2}{\|\Delta \mathbf{Y}_{CO_2e}\|^2}\right) \times 100\% = (1 - sin^2(\varphi)) \times 100\%$$

The same analysis of the two dimensional example can be expanded to n-regions, and can be implemented for any variable of interest. Table 3.1 shows the population-, output- and area-weighted angle calculations for temperature changes (ΔT) and precipitation changes (ΔP). The smaller temperature angle values show that SRM compensates better for regional temperature changes than for regional precipitation changes.

Table 3.1: Angles (φ) calculated using different weighting measures (degrees)

	ΔT	ΔP
Population	3^o	11^o
Output	4^o	23^o
Area	7^o	17^o

Table 3.2 showcases the simplicity of the RCR model. We can compare different SRM levels and analyze their impacts, both for temperature and for precipitation, in terms of

three different social objectives. Policy- and decision-makers can compare different proposals relative to the best case scenario for a given social objective. In the next section we demonstrate how the RCR model can be used to compare sub-optimal policies.

Table 3.2: Percentage of CO_2-driven damages that can be compensated for with SRM for given climate indicator and economic weighting.

		ΔT	ΔP
Best Case	Egalitarian	99	97
	Utilitarian	99	85
	Ecocentric	99	91
Temperature / Utilitarian	Egalitarian	99	-51
	Utilitarian	99	47
	Ecocentric	98	87
Precipitation / Egalitarian	Egalitarian	70	97
	Utilitarian	69	79
	Ecocentric	72	72

The above method provides a parsimonious way to account for regional inequality in the assessment of SRM effectiveness and allows policy and decision makers to examine the linear climate response to different SRM configurations. We found that an SRM scheme optimized to restore population-weighted temperature changes the baseline compensates for 99% of these changes while an SRM scheme optimized for population-weighted precipitation changes compensates for 97% of these changes. Of course, this analysis just shows how inequalities impede an optimal solution. Since you can get the same "angle" with one large residual anomaly in one region and perfect compensations in the others as you get with small residual anomalies in all regions, it is not a good metric for demonstrating diversity of regional preferences. Still, it demonstrates that while inequalities do diminish the potential effectiveness of SRM, they may not be as severe an impediment as it is often assumed.

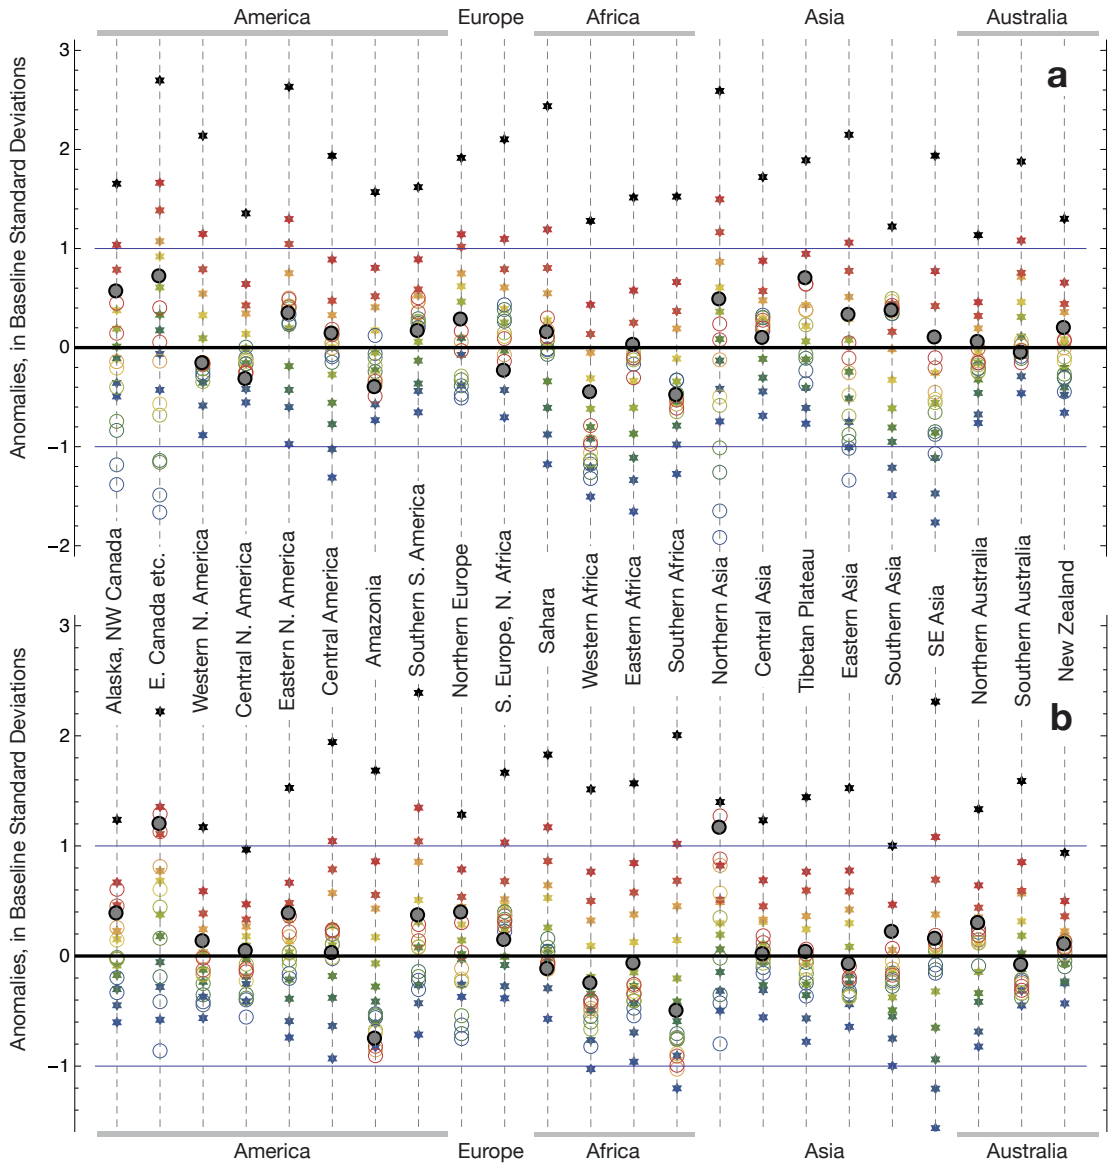

Figure 3.12: Normalized Regional Temperature and Precipitation Anomalies computed as the difference between average values for the decade of the 2020s and the 1990s in units of baseline interannual standard deviations. Stars (*) show mean temperature data for SRM (colors) and no-SRM (black) scenarios. Open circles show mean precipitation data for (a) the Northern Hemisphere summer (JJA), and (b) the Northern Hemisphere winter (DJF). Horizontal lines on each plot show the cut off for one baseline standard deviation. Each colored point represents data from 60 simulations including 6 SRM scenarios. Each black point represents data from 30 no-SRM simulations. The baseline dataset is compiled from 54 simulations.

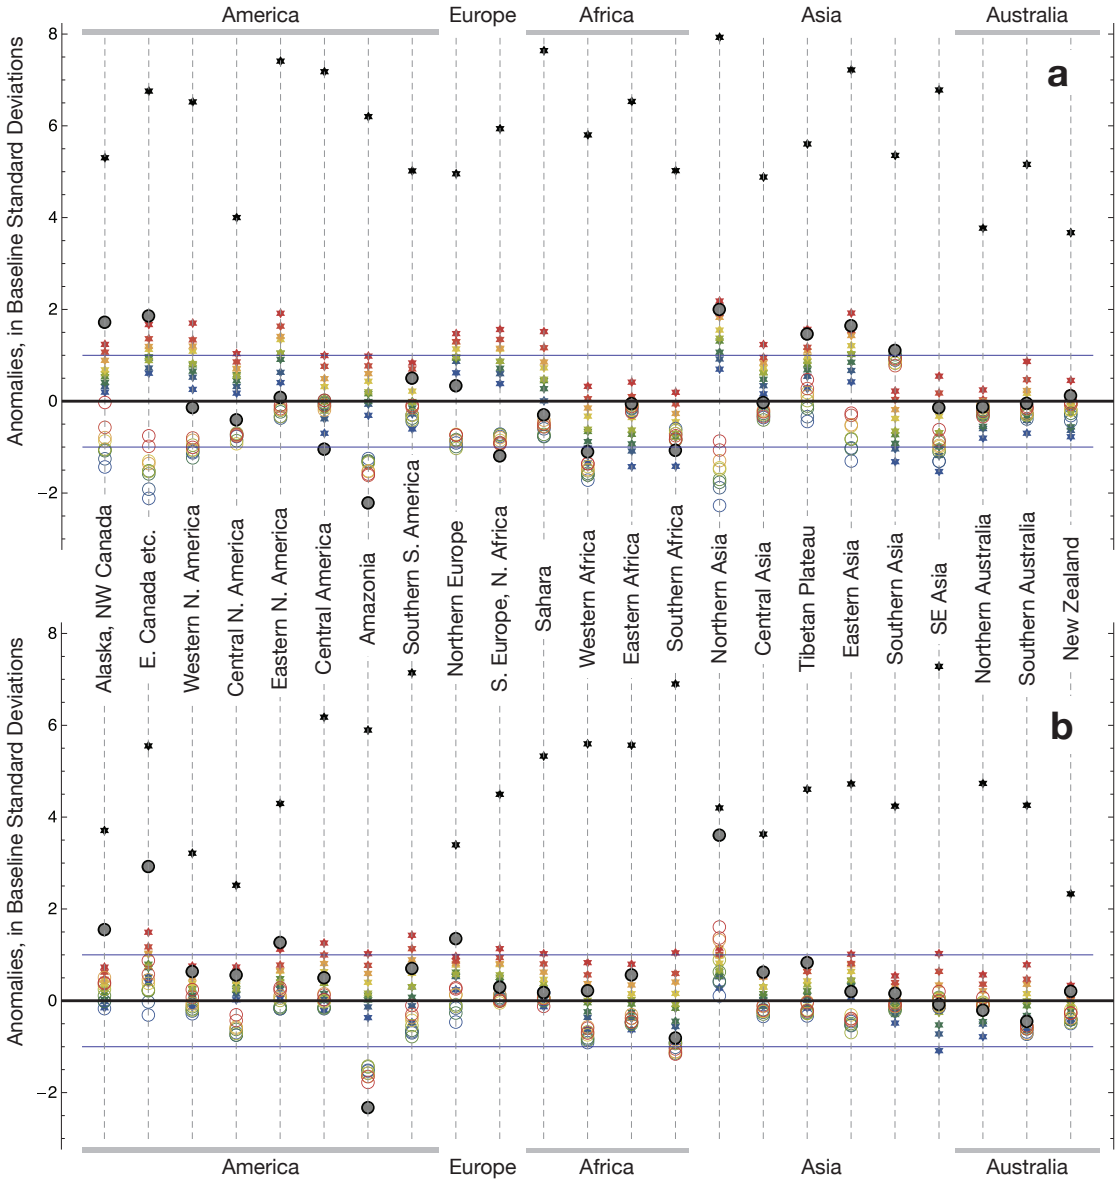

Figure 3.13: Normalized Regional Temperature and Precipitation Anomalies computed as the difference between average values for the decade of the 2070s and the 1990s in units of baseline interannual standard deviations. Stars (*) show mean temperature data for SRM (colors) and no-SRM (black) scenarios. Open circles show mean precipitation data for (a) the Northern Hemisphere summer (JJA), and (b) the Northern Hemisphere winter (DJF). Horizontal lines on each plot show the cut off for one baseline standard deviation. Each colored point represents data from 60 simulations including 6 SRM scenarios. Each black point represents data from 30 no-SRM simulations. The baseline dataset is compiled from 54 simulations.

Chapter 4

Results: Experiment 2

The least controversial justification for learning how to modify the planet's albedo is to provide insurance against the effects of unexpectedly large anthropogenic climate change. A number of modeling studies suggest that SRM can compensate for many of the temperature and precipitation changes associated with global warming, even at the regional level. (Caldeira and Wood, 2008; Lunt et al., 2008; Ricke et al., 2010; Robock et al., 2008) However, these studies have all used only "best estimate" model configurations and parameter sets and so do not account for the uncertainty in the climate system response to SRM arising from uncertain components of the AOGCMs. For this reason, little is known about how SRM's effectiveness and associated inequities may differ in a world with high climate sensitivity in which SRM is most likely to be deployed.

This chapter presents simulations of SRM from a perturbed physics ensemble modelling experiment (as explained on the next page) which have a wide range of temperature responses and climate sensitivities, all of which are consistent with observed recent warming. The analysis shows that the efficacy and distribution of effects of SRM varies with the temperature response of the model. Models that produce more global warming are also generally more sensitive to SRM, so the amount of modification of the Earth's en-

ergy balance needed to meet any given climate stabilization criteria appear to be relatively insensitive to climate sensitivity. However, in more sensitive models, SRM is generally less successful in returning regional climates to their unperturbed states. [1]

4.1 The Perturbed Physics Ensemble

We perform a "perturbed physics" ensemble (PPE) modelling experiment with the climate*prediction*.net (CPDN) version of the HadCM3L AOGCM (Frame et al., 2009; Gordon et al., 2000) to investigate how SRM might behave in models with high climate sensitivity that nevertheless are consistent with recent observed climate change. Like other perturbed physics climate modelling experiments (Murphy et al., 2004; Stainforth et al., 2005), we simulate past and future climate scenarios using a wide range of model parameter combinations with the aim of both reproducing past climate within a specified level of accuracy and also projecting future climates that exhibit a wide range of climate sensitivities.

The standard versions of AOGCMs have generally benefited from considerable tuning – the set of values of model parameters has evolved through a great deal of trial-and-error to one that gives a good simulation of present-day climate. A PPE deliberately "detunes" the model, setting parameters to any value which in principle is physically acceptable, to create diversity. Many of the original 1,550 CPDN models provide a very poor simulation of the recent observed climate. (Rowlands et al., 2011) The goal of our experimental design is to restart models that provide a credible simulation of the past 50 years while maintaining a large diversity in the response in 2050. A number of the choices we make in the design are for pragmatic reasons rather than any desire to apply a formal statistical algorithm,

[1] This chapter is based on a shorter article on results from this experiment, written with Dan Rowlands, M. Granger Morgan, William Ingram and David Keith for submission to *Nature Climate Change*. Dan and I designed the ensemble and Dan generated it. I performed all the data analyses and drafted the article and Granger, William and David contributed to the interpretation and writing. (Ricke et al., 2011)

since we do not seek to interpret the distribution of models in the new ensemble in any probabilistic terms. Several factors were considered in selecting model runs.

First, we fixed the future solar forcing scenario (Solanki and Krivova, 2003), and the future anthropogenic sulphate emissions trajectory. To avoid discontinuities in the solar forcing at the year 2000 we consider only simulations that have a solar forcing very close to the fixed scenario in 2000.

Second, we checked that each model version has a relatively stable base climate. A unique flux adjustment and readjustment approach is taken to spin up of each HadCM3L perturbed physics variant (Frame et al., 2009), which can sometimes fail to eliminate drift in the base control climate of each model. We therefore eliminated model versions in which the initial condition ensemble average of the control simulations exhibited a drift greater than 0.5K/century based on the 1960-2080 portion of the simulation (the first 40 years are generally taken as a spin up phase where the perturbed atmosphere and ocean equilibrate).

Finally, we select model versions through a comparison of the model and observed spatio-temporal pattern of temperature response over the past 50 years. Ensemble members were analyzed using an algorithm for model selection of regional surface temperature data from the cpdn simulation output for model years 1961-2005 (in five year averages) to observational data from the HadCRUT3 (land) (Brohan et al., 2006) and HadSST2 (ocean) (Rayner et al., 2006) datasets at the same spatial and temporal resolution. The algorithm can be thought of as a weighted sum of square differences between the model simulations and observations (which we term r^2, the Mahalanobis distance) (Rowlands et al., 2011). We only included a model for consideration if it had an r^2 less than the 95th percentile of the distribution of r^2 arising from an estimate of internal variability (aleatoric uncertainty) alone. In other words, we reject a model if its r^2 has less than a 5% chance of being generated by internal variability alone. This is akin to saying that we reject the null

hypothesis that differences between a model version and observations are due to internal variability alone

Coupling these constraints with the requirement that a model version must have at least one available restart file in 2000 left us with 238 model versions (with 557 associated restart files), denoted by the open gray triangles under the dashed gray horizontal line in Figure 4.1. Small grey triangles show simulations without restart files available in 2000.

To select a subset of the 238 models for inclusion in the new ensemble that ensured a wide range of responses in the future, the models were binned by forecast warming in 2050 into 10 equally spaced bins spanning the range of responses. In each bin, the model version with the lowest r^2 was automatically included, along with 4 others sampled probabilistically, ensuring there are no duplicates. In the two highest response bins there are less than 5 model versions that met the selection criteria, and hence our selection yielded only 43 model versions.

Figure 4.1 shows the goodness of fit between model and observation differences against simulated warming in 2050 with our forty-three member PPE ensemble in color. The color code for the points indicates the model's calculated climate sensitivity from corresponding equilibrium slab ocean simulations, which is correlated with temperature response.

For each model version selected there are n possible restart files. We selected 10 random numbers with replacement from 1 to n, which then gives a set of 10 (often non-unique) restart files for each model version. For each of the 430 restart files, we attached an initial condition perturbation (a temperature perturbation in a single model cell) selected as a random number from 0 to 0.2. Hence for each model version we start from a number of different restart files and then perturb off each of these in creating an initial condition ensemble. For each of the 4 SRM scenarios used for the analysis in this paper, this 430 member ensemble was repeated, giving a total of 1720 model simulations.

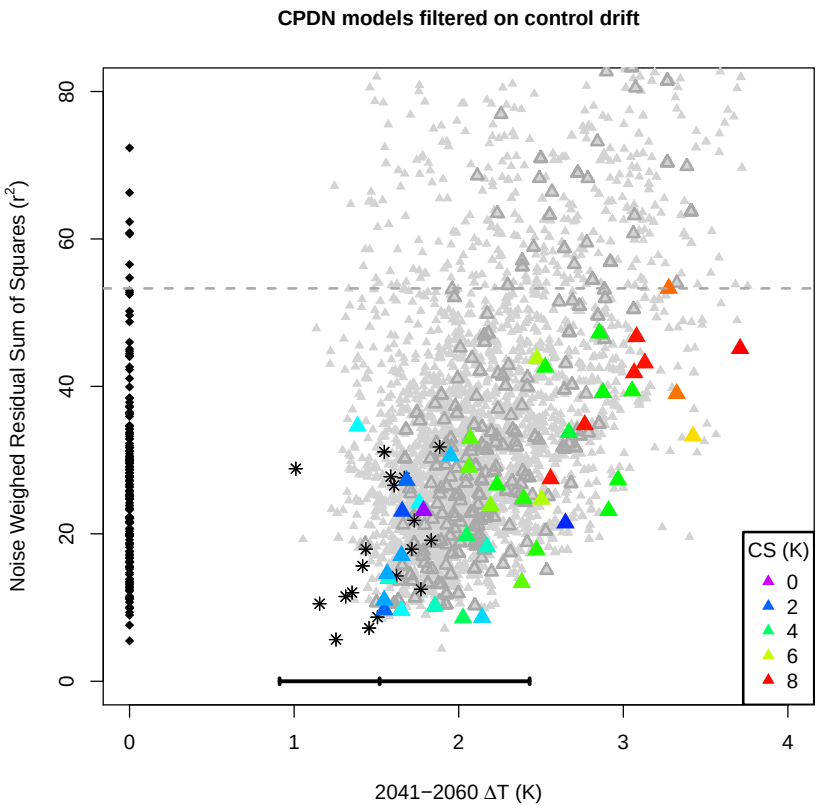

Figure 4.1: Residual sum of squares from observations against 2041-2060 forecast warming relative to 1981-2000, for the climate*prediction*.netensemble under the SRESA1B scenario. Each point refers to a particular model version (physics combination) and so is the average over initial condition members. Small grey triangles indicate all simulations, those with large grey triangles indicating those available for sampling. Coloured triangles indicate the models sampled, shaded by the model's climate sensitivity. For comparison the performnce of models used in the IPCC fourth assessment are denoted by black stars. The black bar at the bottom of the plot indicates the CMIP3 mean and range. Black dots on the left indicate the range in the residual arising from natural variability alone as defined by pre-industrial control simulations from the CMIP3 archive, with the horizontal dotted line indicating the 95th percentile of this distribution.

4.2 SRM Scenarios

Figure 4.2 shows a plot of the scenarios over the length of the simulations. SRM activities were implemented in the model as precribed values of stratospheric aerosol optical depth

(AOD), as described in Chapter 2. A baseline SRM scenario (medium-SRM) was formulated using the results from the standard physics experiment (Ricke et al., 2010) in which 135 SRM scenarios were formulated, designed to offset the net forcings associated with long-lived greenhouse gases, tropospheric sulphur aerosols and tropospheric ozone; and spanning the uncertainties associated with these anthropogenic forcings. The two scenarios from that experiment which best stabilized global surface air temperature according to a least-squares fit analysis were averaged. The high- SRM and low-SRM scenarios are the same as the baseline scenario except for the addition (0.075) or subtraction (0.015) of a constant amount of optical depth. These adjustments were calibrated using the output from the standard physics experiment data to bracket a full range of potential ideal responses to SRM for all of the 23 regions analyzed.

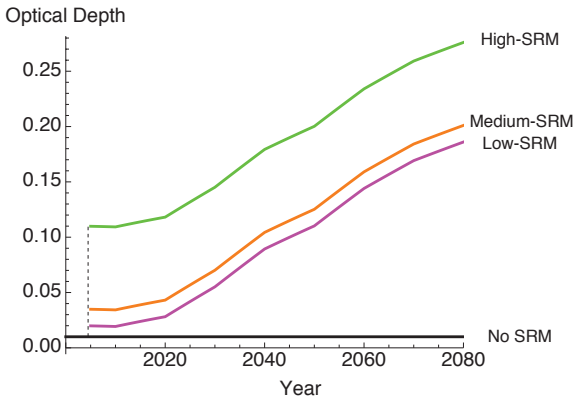

Figure 4.2: Three SRM Forcing scenarios, and one no-SRM scenario implemented as set values of stratospheric aerosol optical depth.

4.3 Global Results

Figure 4.3 shows five-year running-mean global-mean surface air temperatures and precipitation rates for the no-SRM, low-SRM and high-SRM scenarios. SRM with strato-

spheric aerosols cannot simultaneously compensate for the impacts of rising greenhouse gases on both temperatures and the hydrological cycle, because while SRM and greenhouse gases have opposite effects on temperature (and thus on the temperature-forced response of precipitation), both have the direct effect of reducing precipitation. (Allen and Ingram, 2002; Bala et al., 2008) Under the no-SRM scenario, global temperatures and precipitation increased in all of the models. While results vary across models, both high- and low-SRM temperature is relatively stable after 2020 and precipitation falls.

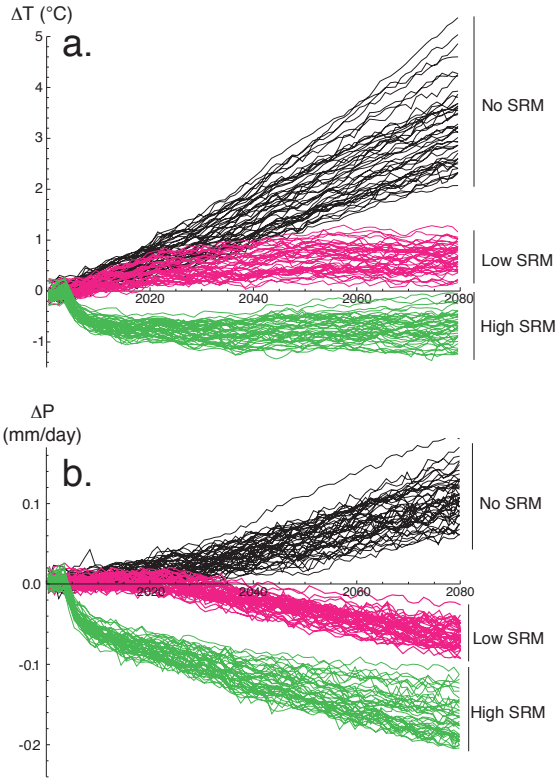

Figure 4.3: Time series of temperature and precipitation of the no-SRM, low-SRM and high-SRM scenarios examined, with initial condition sub-ensembles averaged for each of the forty PPE model configurations analyzed. (**a**) Five-year average global mean near-surface (1.5 m) air temperature, and (**b**) five-year average global mean precipitation rate, all displayed over the length of the 80 model-year simulations.

4.4 Measures of Regional SRM Effectiveness

To analyze the regional effects of different levels of SRM we examined mean temperature and precipitation anomalies over land in the 22 "Giorgi regions" (Giorgi and Francisco, 2000) plus New Zealand as described in Chapter 2. Figure 4.4 shows an example of how regional responses to greenhouse gas and SRM forcings vary between models. It shows decadal-mean temperature and precipitation changes between 2000 and 2050, normalized by the ensemble-mean inter-annual variability of their control (unperturbed by greenhouse gases or SRM) climates, for two regions and for both the standard settings of model physics parameters (ΔT_{2050}=2.1°C) and our ensemble's highest temperature response model configuration (ΔT_{2050}=4.1°C).

With both model configurations, Eastern North America gets warmer and wetter under A1B, while the Mediterranean gets warmer and drier. When SRM is used, both regions move back towards their baseline climate states in both models and in the standard physics model, with the right amount of SRM, the regions could return almost exactly to the 2000 baseline for both temperature and precipitation. The best SRM can do in the high response climate, however, is to bring each region back within one standard deviations of the origin. It is also important to note that, unlike the Eastern North America and the Mediterranean, for some regions it is not possible to return temperature and precipitation to their baseline values even in the standard physics model. (See Appendix E.)

The ensemble design, which samples evenly among models of different temperature responses, is well suited for an analysis that examines the relationship between various regional measures of SRM efficacy and the overall global temperature response or climate sensitivity of the model. The following measures are designed to gauge how well SRM could work, in compensating for the effects of global warming in a single region and how well it works to restore regional climates overall and relative to no-SRM, once a particular

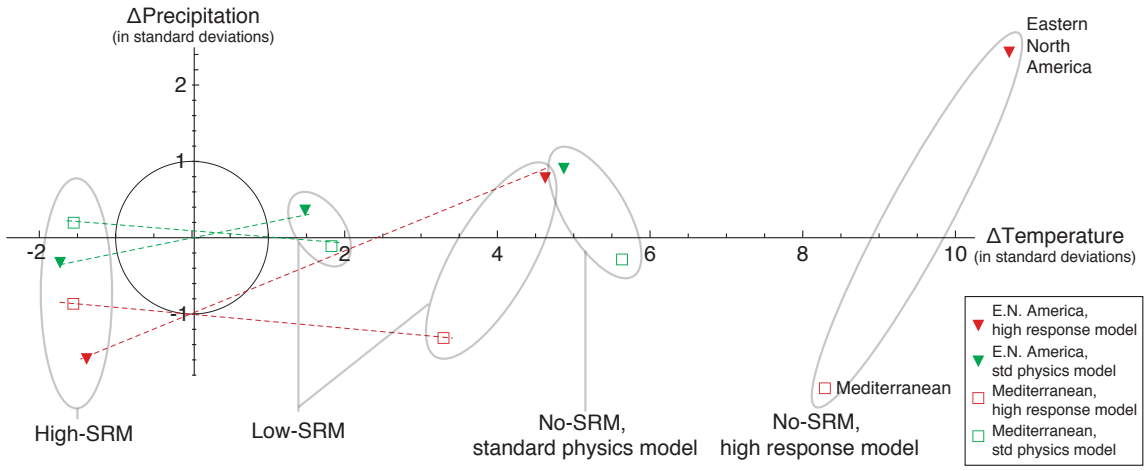

Figure 4.4: Example of how regional responses to greenhouse gas and SRM forcings vary between models. The green points show the no-SRM, low SRM and high SRM responses for the standard settings of model physics parameters (ΔT_{2050}=2.1°C), while the red points show the responses our ensemble's highest temperature response model configuration ΔT_{2050}=4.1°C). Changes in temperature and precipitation are calculated as the ten-year average 2050 minus 2000, divided by the interannual variability of the unperturbed control climate (calculated with simulations that used identical parameter combinations but no anthropogenic forcings of any kind). Dashed lines indicate the linear trajectory as SRM is increased or decreased.

level of SRM is set.

First, we consider OD*, the amount of optical depth modification that returns a given regional climate closest to its baseline state (the origin in Figure 4.4) in units of standard deviations, to measure the diversity of regional preferences for the amount of SRM. Second, we consider regional anomalies, i.e., the regional temperature, precipitation or net temperature and precipitation changes, associated with a set amount of SRM to gauge likely regional satisfaction with and potential inequities in SRM's compensatory power relative to their baseline climate state. The amount of SRM is "democratically selected": simply the mean value of OD* for all regions, either unweighted (each region gets a vote), or weighted by regional population or economic output (each person or dollar gets a vote). (See Chapter 2 for information on the socioeconomic datasets used.) Finally, we consider

regional anomalies relative to no those associated with no SRM to gauge likely regional satisfaction with and potential inequities in SRM's compensatory power relative to "doing nothing" to counteract the effects of the climate changes from rising greenhouse gases. (See Appendix F for more precise definitions of the metrics and their derivation.)

These metrics are a tool to explicitly account for the precipitation and temperature tradeoffs that occur when attempting to stabilize regional climate changes using SRM. Precipitation and temperature changes are only two, albeit very important, of the many variables likely to have climate related impacts. They are not intended as definitive normative measures of regional impacts or preferences. For example, when temperature rises in a region, an increase in precipitation is required to maintain soil moisture. Any region would likely prefer a shift in their mean climate state to an increase of half a standard deviation in temperature with a half standard deviation more precipitation rather than to an increase of half a standard deviation in temperature with half a standard deviation less precipitation. However, incremental changes in the amount of SRM never result in such a shift in precipitation space only. These metrics are specific to assessing SRM efficacy, but are obviously less informative with respect to a number of other alternative policies for dealing with climate change.

In reality, the relative importance of temperature and precipitation changes on impacts is likely to be region-specific. A recent paper by Lobell et al. (2011) which estimated the effects of temperature and precipitation changes on crop yields found that a change of one standard deviation in precipitation often had approximately the same effect as a change in one standard deviation in temperature, but that this varied considerably by country and by crop. For example, they found that for the United States, a one standard deviation decrease in precipitation had a larger negative impact than a one standard deviation increase in temperature on crop yields of corn and soy, but a smaller negative impact than temperature on yields of wheat. Likewise, in China, a one standard deviation

decrease in precipitation had a larger negative impact than a one standard deviation increase in temperature on crop yields of corn and soy, but a one standard deviation change in precipitation was predicted to increase wheat yields.

4.5 Efficacy and Equity Across Regions

Figures 4.5-4.8 show the ten-year mean values of various efficacy measures by model temperature response for decades averaged around 2030, 2050 and 2070. As greenhouse gas concentrations rise over the length of the simulations, increased SRM is required to compensate and as such, OD* rises across all model configurations. As illustrated in Fig. 4.5, mean regional preferences for the amount of optical depth modification are fairly insensitive to model temperature response. This result is robust to different socioeconomic weightings if weighted by regional population or economic output. The standard deviation of regional preferences, shown in Figure 4.6, decreases with model temperature response. This should be expected physically as the smaller differences between the different levels of SRM have more impact on higher sensitivity climates, just as smaller changes in greenhouse gases do.

Figure 4.7 shows that mean regional anomalies, independent of weighting scheme, rise with model temperature response, but these results are only statistically significant ($p<0.05$) for the anomalies (all weightings) in 2070, the economic output-weighted anomalies in 2050 and the population-weighted anomalies in 2030. (See Table 4.1 for regression coefficients and p-values for all best-fit lines presented in the figures.) As a proxy for regional satisfaction with SRM, this result implies that satisfaction with SRM may be overall lower in higher sensitivity worlds. The fact that SRM may be a more imperfect substitute for mitigation the higher the temperature response of the global climate is to greenhouse gas forcings is noteworthy, as these are the situations in which SRM is particularly likely

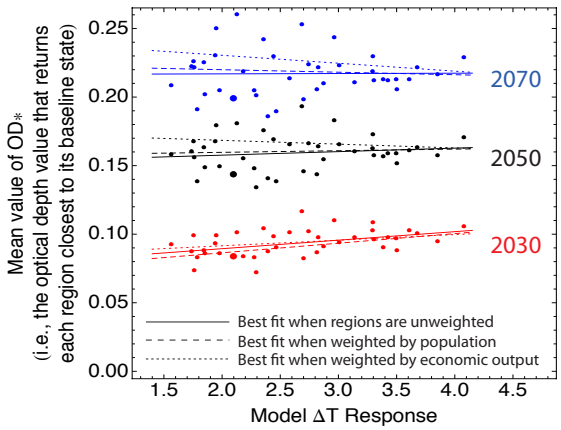

Figure 4.5: Mean regional values of OD*, the amount optical depth modification that returns a given regional climate closest to its baseline state (i.e. to the origin in Figure 4.4), plotted as a function of model temperature response for decadal intervals centered on 2030, 2050 and 2070. Colored points report the results for the sampled model runs that use a specific set of parameter values (see colored points in Figure 4.1) in which equal weight has been given to each of the 23 regions. Solid lines show best fit to these points. Dashed and dotted lines show best fits to points (not shown) that result if each geographic region is weighted by its economic output (dotted) or by its population (dashed).

to be deployed.

The standard deviation of regional anomalies also rises with model temperature response, suggesting inequities in satisfaction associated with an SRM substitution would be greater in higher sensitivity worlds. Amplified regional drying in high-response worlds drives the overall higher regional anomalies in Fig. 4.7a. Figures 4.7c and 4.7d show that on average across the ensemble, regions are a bit warmer and more dry with an SRM-modified climate. There is no significant relationship between model temperature response and the magnitude of regional temperature anomalies.

Figure 4.8 shows that the mean ratio between regional anomalies at OD* and their anomalies without any SRM falls with model temperature response and falls over the length of the simulations, as does the standard deviation of this ratio. The best-fits for this measure are all statistically significant ($p<0.05$). Viewed in this way, SRM is more

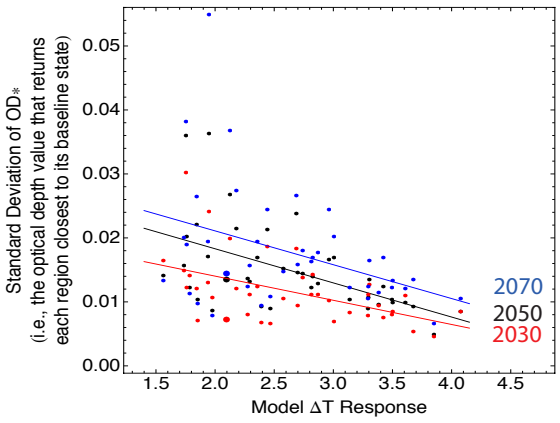

Figure 4.6: The standard deviation of regional values of OD*, the amount optical depth modification that returns a given regional climate closest to its baseline state (i.e. to the origin in Figure 4.4), plotted as a function of model temperature response for decadal intervals centered on 2030, 2050 and 2070. Display conventions are the same as those in Figure 4.5.

effective at reducing the amount of risk from climate change in high climate sensitivity worlds. The higher relative effectiveness of SRM in high sensitivity worlds is especially apparent if you consider the extent to which SRM can stabilize rate of change, as illustrated in Figures 4.9 and 4.10. On average, regional rates of temperature change are twice as high in the highest sensitivity models than in the lowest sensitivity models, whereas with SRM applied, the rates of temperature change are completely uncorrelated the model's climate sensitivity and rates of precipitation change are only slightly higher in the higher sensitivity models.

4.6 The Entrainment Coefficient

The process by which the air inside a cloud and the air around it mix together (exchanging heat, moisture and liquid water) is called entrainment. It is a sub-grid process which is parameterized in HadCM3L as a part of the convective cloud scheme. (Gregory and Rown-

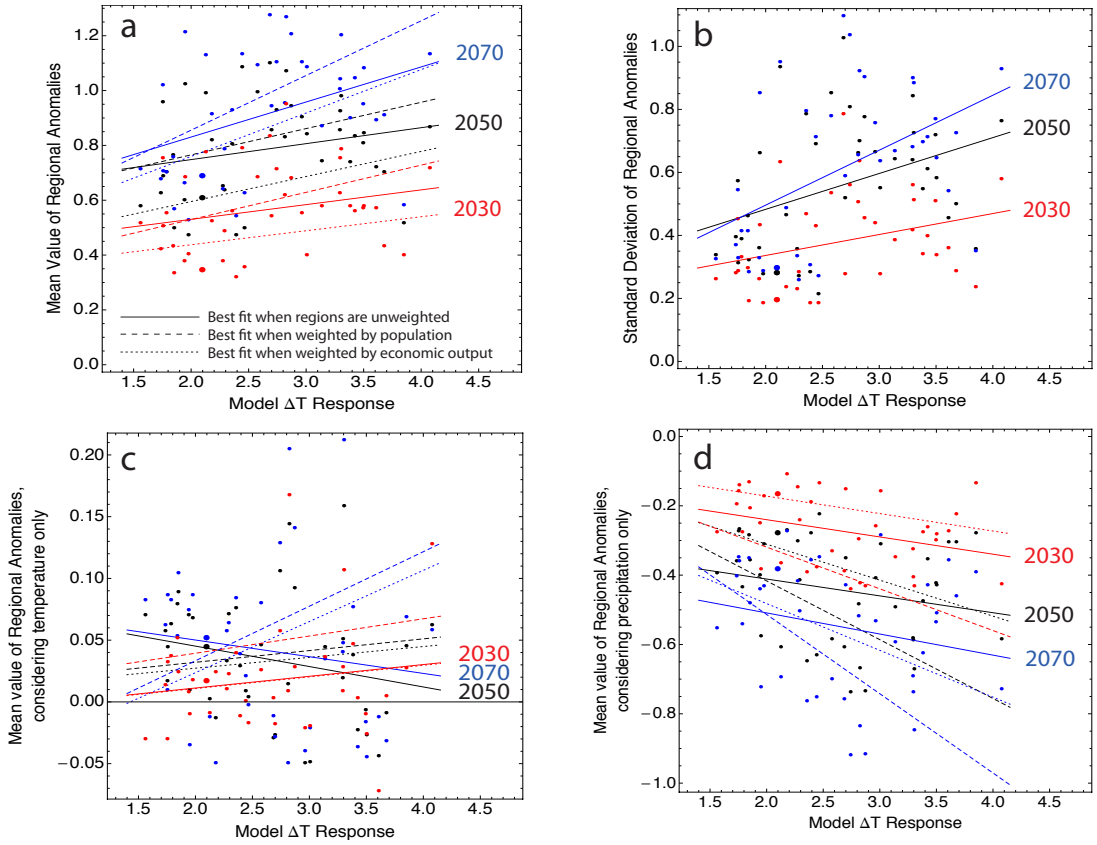

Figure 4.7: Statistics on regional anomalies, as measured in temperature and precipitation standard deviation space. Plots (**a**) and (**b**) show the mean and standard deviation of the regional combined temperature-precipitation anomalies when the level of SRM is set at mean-OD* in units of control climate inter-annual standard deviations, plotted as a function of model temperature response for decadal intervals centered on 2030, 2050 and 2070. Separate indications of the mean value of regional anomalies in temperature (**c**) and precipitation (**d**). Display conventions are the same as those in Figure 4.5.

tree, 1990) The important influence of the entrainment coefficient parameter in HadCM3L on the CPDN ensemble results is well documented. (Knight et al., 2007; Stainforth et al., 2005) Likewise, in the results from this experiment, differences in the entrainment coefficient generally explain the most variance in the simulation output.

The analysis in the previous section was predicated on the hypothesis that some relationship exists between the climate response to GHG forcings and its response to SRM.

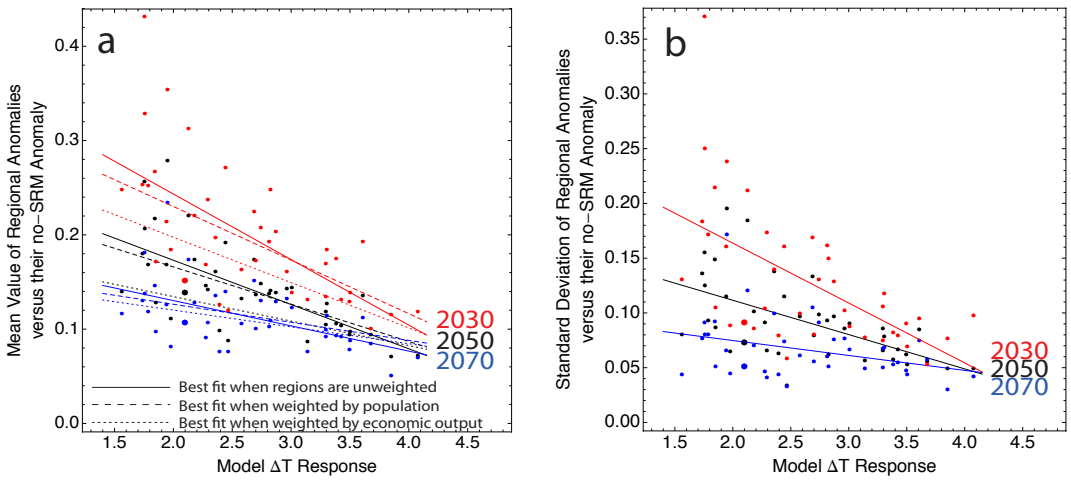

Figure 4.8: The mean and standard deviation of the ratio of regional temperature-precipitation anomalies versus the anomalies with no SRM when the level of SRM is set at mean-OD* in units of control climate inter-annual standard deviations, plotted as a function of model temperature response for decadal intervals centered on 2030, 2050 and 2070. Display conventions are the same as those in Figure 4.5.

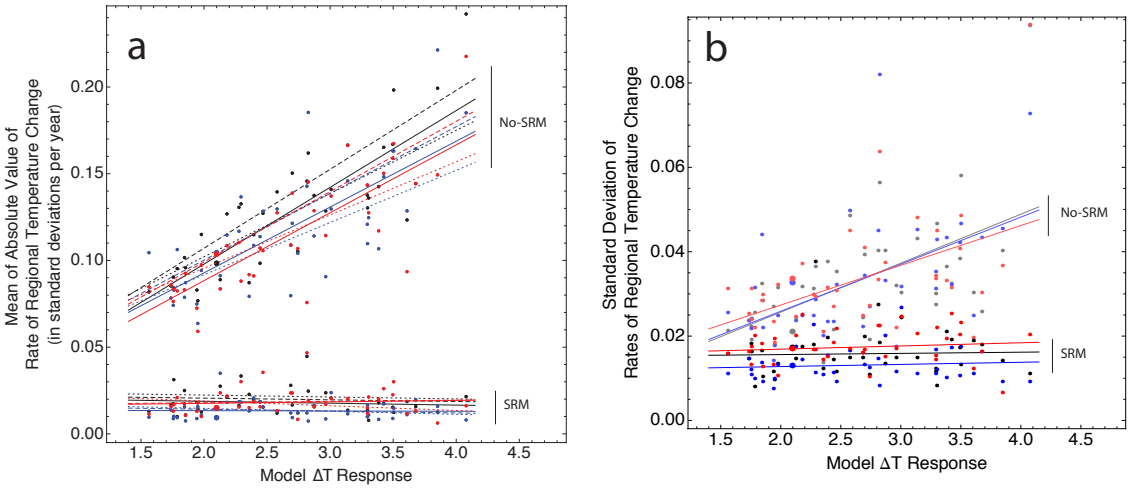

Figure 4.9: Regional rate of temperature change. (a) shows the mean value of the absolute values of regional rate of change for temperature for for SRM at $\overline{OD*}$ and no-SRM scenarios for decadal intervals centered on 2030 (red), 2050 (black) and 2070 (blue) in units of standard deviations per year. (b) shows the standard deviation of the 23 regional rates of change.

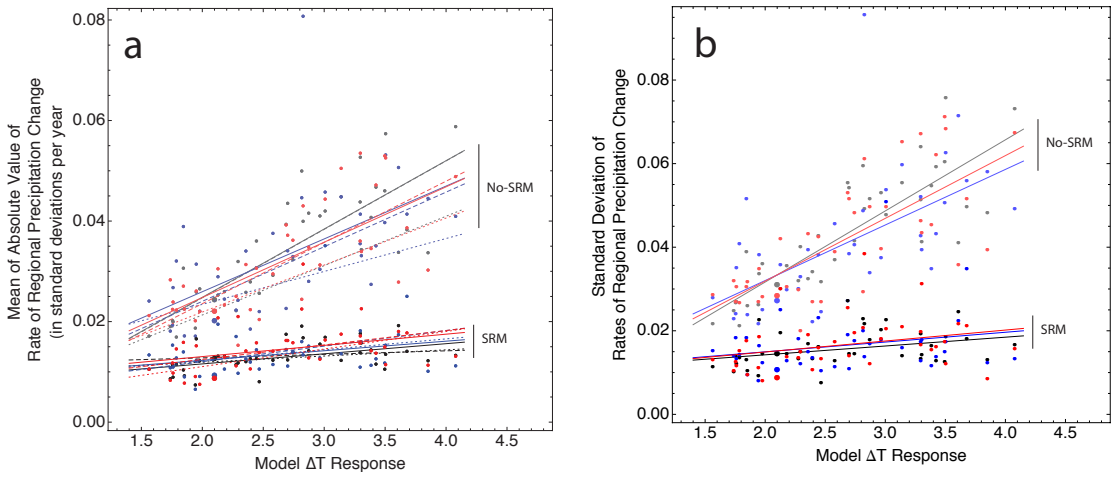

Figure 4.10: Regional rate of precipitation change. (a) shows the mean value of the absolute values of regional rate of change for precipitation for for SRM at $\overline{OD*}$ and no-SRM scenarios for decadal intervals centered on 2030 (red), 2050 (black) and 2070 (blue) in units of standard deviations per year. (b) shows the standard deviation of the 23 regional rates of change.

Significant correlations do exist between SRM efficacy measures and the model forecast warming (or climate sensitivity). In order to take a more agnostic approach to identifying the source of the variation in PPE model response to these forcings, I used regression trees to split data into variance-minimizing groups, including model forecast warming as an explanatory variable. (Breiman et al., 1984) The usefulness of this approach is somewhat limited by the relatively small PPE ensemble size, and conservative criteria were set for splitting data (in particular, the minimum number of observations at a node in order for it to be eligible for a split was 10 of the 43 total observations). Nonetheless, the trees are helpful in demonstrating the powerful influence of the entrainment coefficient. Figures 4.11 and 4.12 show the pruned regression trees for OD-star and mean regional anomaly in 2050. (Testing forecast warming's explanatory power in regard to the relative efficacy measure is not particularly illuminating because 2050 forecast warming is the denominator in that ratio.)

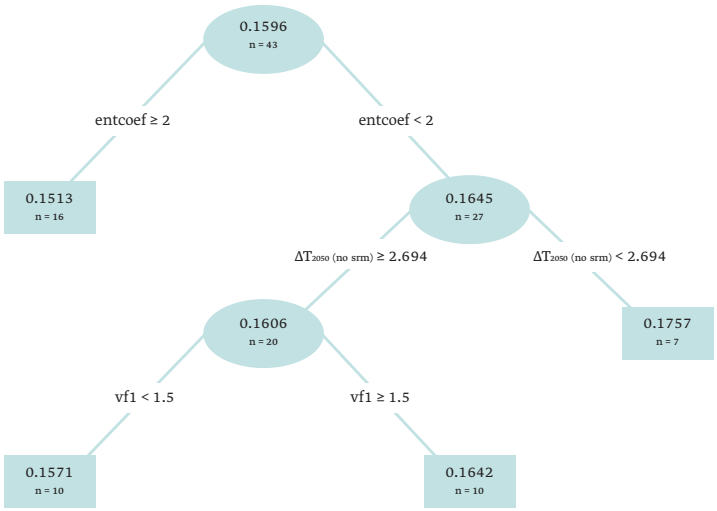

Figure 4.11: Regression Tree for OD* in 2050

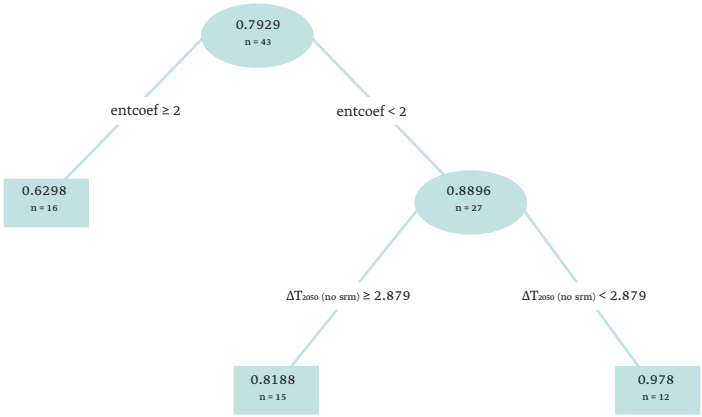

Figure 4.12: Regression Tree for Mean Regional Anomaly in 2050

Splitting the data into subsets that have entrainment coefficients greater than or less than two explains approximately 50% of the variation in OD-star and 26% of the variation in the mean regional anomaly at OD-star (as opposed to the R^2 values of 2% and 16% for a least-squares regression on forecast warming alone, as presented in the previous section).

These results do not invalidate the results above. The entrainment coefficient is a legitimate source of uncertainty in the model and the models are constrained by observations. The dominant influence of this parameter, however, is worth noting.

4.7 Maps of Residual Changes in SRM-Modified Climates

Figures 4.13-4.16, which show seasonal temperature and precipitation maps for the standard physics model and the four highest and lowest temperature response models. The temperature maps show the differences between 2050 and 1990 under the high SRM scenario because it is the scenario that best minimized global-mean temperature changes (without any extrapolation) on average across the PPE. Likewise, precipitation maps show the differences between 2050 and 1990 under the medium SRM scenario because it is the scenario that best minimized global-mean precipitation changes. (The global mean residual changes are shown in the bottom left corner of each map.) By presenting temperature and precipitation changes for the scenario that minimizes mean changes for each indicator best, the colored areas on the map primarily show the residual changes associated with using SRM to compensate for GHGs.

The most extreme residual temperature changes are generally concentrated in higher latitudes, while the most extreme residual precipitation changes (measured in % difference) are concentrated in equatorial and midlatitudes. Beyond these general features, there is considerable geographic diversity between models. While there are significant trends in how the effectiveness of SRM at the regional level varies with climate sensitivity in general, the results in particular regions often do not exhibit a clear trend. Understanding the physical significance of various parametric sensitivities at the regional level is a topic that merits further analysis in future work (see Chapter 6).

4.8 Discussion

Some general conclusions from the standard physics experiment and other SRM modeling exercises are robust to parametric uncertainty: SRM-modified regional climates generally migrate away from their baseline climate state as simulations progress, reducing the potential effectiveness of SRM as a substitute for mitigation the longer it is used in that capacity. In addition, the differences in regional preferences increase the more/longer SRM is used and inequities in the effectiveness of various "democratic" levels of SRM also grow.

There are other important risks and problems associated with SRM that may vary with global temperature response. For example, there are increased risks of stratospheric ozone destruction associated with most proposed forms of stratospheric SRM, (Tilmes et al., 2009) and this problem could be compounded by stratospheric dynamics changing in a warming world. Cooler and wetter conditions in the lower stratosphere increase the conversion of inorganic chlorine to free radical form, resulting in increased potential for ozone destruction. (Kirk-Davidoff et al., 1999) Changes to stratospheric temperatures and humidity could well vary with climate sensitivity, either amplifying or dampening the chain of effects leading to increased ozone destruction. Such a problem would be worth investigating in a model with a better-resolved stratosphere than HadCM3L's.

Results from this experiment demonstrate how model assessments of SRM that use the best estimate parameter sets and so moderate climate sensitivity, may be ignoring some of important contingencies associated with implementing SRM in reality. A primary motivation for studying SRM via the injection of aerosols in the stratosphere is to evaluate its potential effectiveness as "insurance" in the case of higher-than-expected climate response to global warming. We find that this is precisely when SRM appears to be least effective in returning regional climates to their baseline states. On the other hand,

given the very high regional temperature anomalies associated with rising greenhouse gas concentrations in high sensitivity models, it is also where SRM is most effective relative to a no-SRM alternative.

Figures 4.9 and 4.10 suggest that it should be possible to reduce the rate of warming and rate of precipitation changes at the regional level using SRM, with comparable effectiveness under high and low climate sensitivity. Because the mean regional rates of change with SRM are the same at all points in the simulation, the ability of SRM to reduce rates of change in the face of high climate sensitivity does not appear to be a strong function of the decade in which it is employed (or the inter-regional weighting scheme that is employed). All regions would probably find at least some SRM to be preferable to continued rapid change, but is unlikely that all regions will be comparably satisfied with their local outcomes, and many regions may find the result increasingly unsatisfactory over time. Nonetheless, this quality of SRM to reduce rates of regional climate change, even in the face very high climate sensitivity, speaks to the potential power of SRM to make climate change more adaptable.

Table 4.1: Regression coefficients and p-values.

	2030			2050			2070		
Best-Fit Lines in:	α	β	p-value (β)	α	β	p-value (β)	α	β	p-value (β)
Figure 4.5 - unweighted	0.0770	0.0062	0.0041 ***	0.1525	0.0026	0.4069	0.2165	0.0002	0.9521
Figure 4.5 - pop. weighted	0.0724	0.0070	0.0014 **	0.1574	0.0012	0.7175	0.2236	-0.0018	0.6667
Figure 4.5 - econ. weighted	0.0834	0.0041	0.1120	0.1740	-0.0028	0.4814	0.2420	-0.0058	0.2759
Figure 4.6	0.0216	-0.0038	0.0009 ***	0.0290	-0.0053	0.0004 ***	0.0317	-0.0053	0.0128 *
Figure 4.7a - unweighted	0.4228	0.0537	0.1445	0.6318	0.0582	0.1640	0.5747	0.1279	0.0112 *
Figure 4.7a - pop. weighted	0.3310	0.0993	0.0331 *	0.5737	0.0959	0.0521	0.4568	0.1994	0.0018 **
Figure 4.7a - econ. weighted	0.3355	0.0511	0.0888	0.4124	0.0914	0.0152 *	0.4431	0.1583	0.0026 **
Figure 4.7b	0.2033	0.0665	0.0524	0.2551	0.1138	0.0170 *	0.1474	0.1744	0.0016 **
Figure 4.7c - unweighted	-0.0078	0.0095	0.3519	0.0782	-0.0165	0.1559	0.0772	-0.0135	0.3807
Figure 4.7c - pop. weighted	0.0115	0.0139	0.2263	0.0131	0.0095	0.4686	-0.0550	0.0442	0.0428 *
Figure 4.7c - econ. weighted	-0.0081	0.0094	0.3708	0.0097	0.0088	0.4909	-0.0592	0.0414	0.0528
Figure 4.7d - unweighted	-0.1405	-0.0498	0.0373 *	-0.3141	-0.0486	0.1714	-0.3872	-0.0609	0.1537
Figure 4.7d - pop. weighted	-0.0786	-0.1203	0.0068 **	-0.0774	-0.1696	0.0028 **	-0.0563	-0.2286	0.0017 **
Figure 4.7d - econ. weighted	-0.0715	-0.0503	0.0480 *	-0.1031	-0.1038	0.0027 **	-0.2121	-0.1350	0.0021 **
Figure 4.8a - unweighted	0.3823	-0.0695	0.0000 ***	0.2670	-0.0470	0.0000 ***	0.1839	-0.0268	0.0003 ***
Figure 4.8a - pop. weighted	0.3437	-0.0568	0.0002 ***	0.2450	-0.0395	0.0003 ***	0.1647	-0.0191	0.0303 *
Figure 4.8a - econ. weighted	0.2936	-0.0481	0.0001 ***	0.1865	-0.0260	0.0016 **	0.1558	-0.0178	0.0256 *
Figure 4.8b	0.2731	-0.0547	0.0001 ***	0.1742	-0.0313	0.0001 ***	0.1021	-0.0136	0.0268 *
Figure 4.9a- No-SRM - unweighted	0.0100	0.0392	0.0000 ***	0.0094	0.0443	0.0000 ***	0.0169	0.0380	0.0000 ***
Figure 4.9a- No-SRM - pop. weighted	0.0181	0.0405	0.0000 ***	0.0161	0.0455	0.0000 ***	0.0232	0.0385	0.0000 ***
Figure 4.9a- No-SRM - econ. weighted	0.0342	0.0307	0.0000 ***	0.0292	0.0365	0.0000 ***	0.0315	0.0301	0.0000 ***
Figure 4.9a- - unweighted	0.0442	-0.0075	0.2518	0.0252	-0.0022	0.2990	0.0135	0.0015	0.3453
Figure 4.9a- SRM - pop. weighted	0.0478	-0.0079	0.2777	0.0259	-0.0023	0.3629	0.0183	-0.0002	0.9141
Figure 4.9a- SRM - econ. weighted	0.0477	-0.0075	0.2619	0.0296	-0.0030	0.2602	0.0203	-0.0008	0.6946
Figure 4.9b - no-SRM	0.0085	0.0094	0.0017 **	0.0025	0.0116	0.0001 ***	0.0029	0.0114	0.0002 ***
Figure 4.9b -SRM	0.0195	-0.0016	0.3052	0.0152	-0.0009	0.3764	0.0112	0.0003	0.7023
Figure 4.10a- No-SRM - unweighted	0.0028	0.0110	0.0000 ***	-0.0024	0.0136	0.0000 ***	0.0049	0.0105	0.0001 ***
Figure 4.10a- No-SRM - pop. weighted	-0.0008	0.0122	0.0000 ***	-0.0025	0.0136	0.0000 ***	0.0023	0.0109	0.0004 ***
Figure 4.10a- No-SRM - econ. weighted	0.0032	0.0094	0.0000 ***	0.0017	0.0098	0.0000 ***	0.0103	0.0065	0.0096 **
Figure 4.10a- - unweighted	0.0168	0.0006	0.6644	0.0101	0.0025	0.0243 *	0.0108	0.0025	0.0273 *
Figure 4.10a- - pop. weighted	0.0150	0.0011	0.4450	0.0114	0.0023	0.0819	0.0072	0.0042	0.0055 **
Figure 4.10a- - econ. weighted	0.0140	0.0010	0.5063	0.0096	0.0026	0.0177 *	0.0112	0.0025	0.0660
Figure 4.10b	-0.0010	0.0108	0.0000 ***	-0.0015	0.0113	0.0000 ***	0.0033	0.0086	0.0002 ***
Figure 4.10b	0.0138	0.0008	0.6165	0.0077	0.0021	0.0303 *	0.0096	0.0015	0.1078

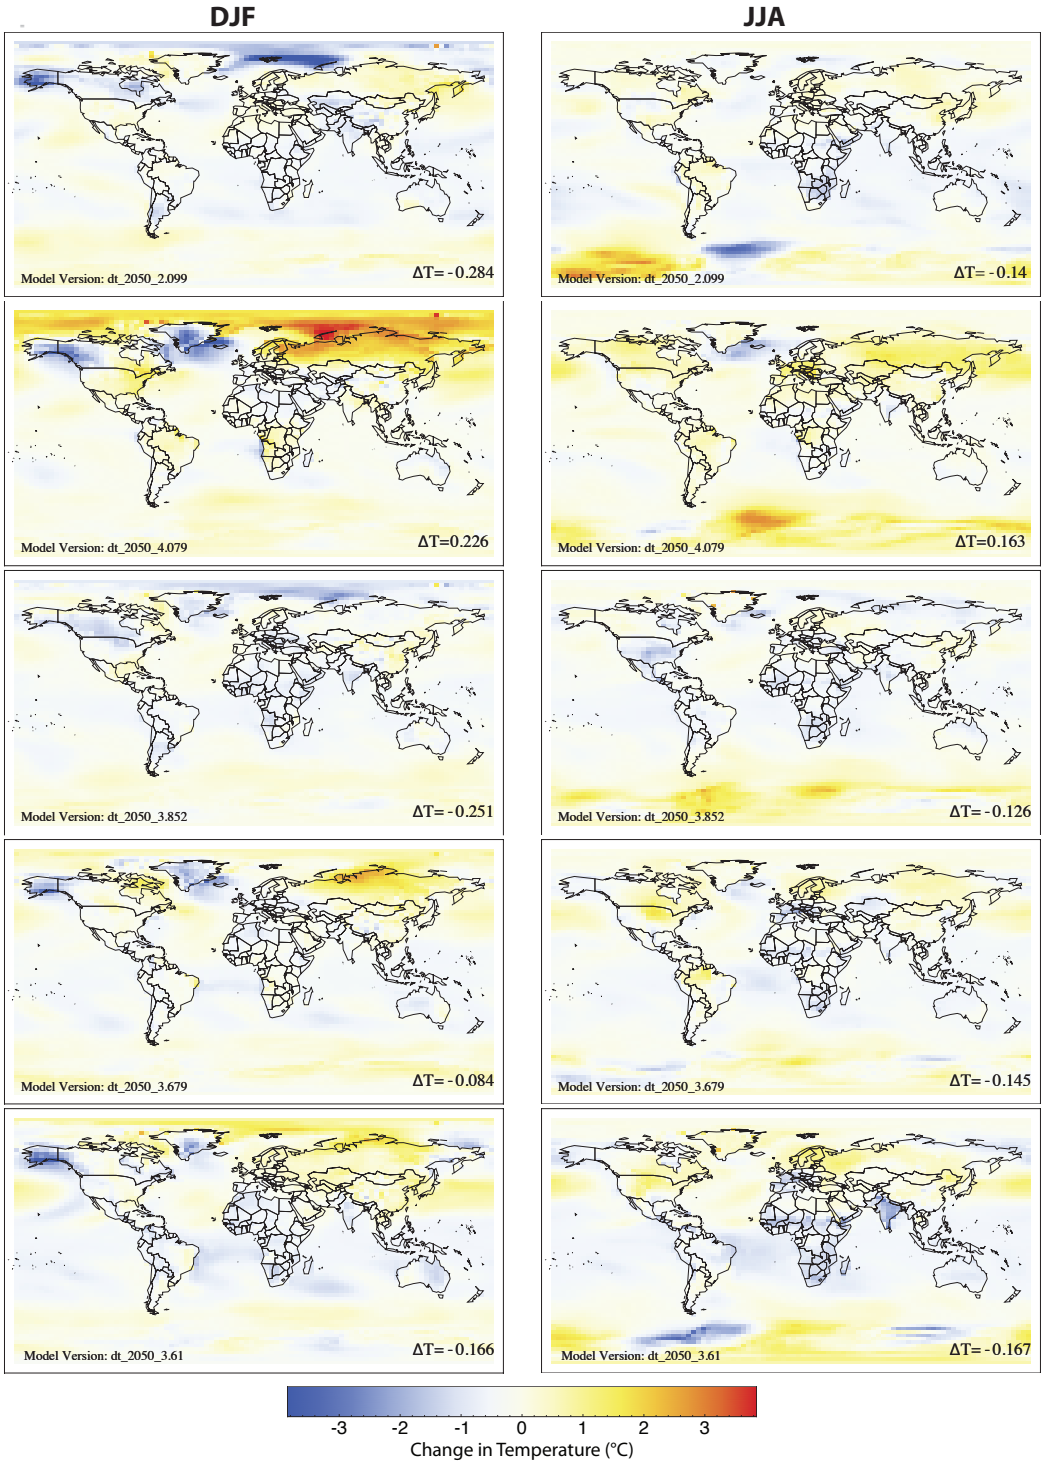

Figure 4.13: 20-Year Seasonal-Mean Near-Surface (1.5 m) Air Temperature Anomaly in °C (2050-1990) for high-SRM scenario. Northern Hemisphere Winter (DJF) is on the left and Northern Hemisphere Summer (JJA) is on the right Maps for simulations using the standard physics model (top), followed by the four highest temperature response models. Mean of three-member initial condition ensembles.

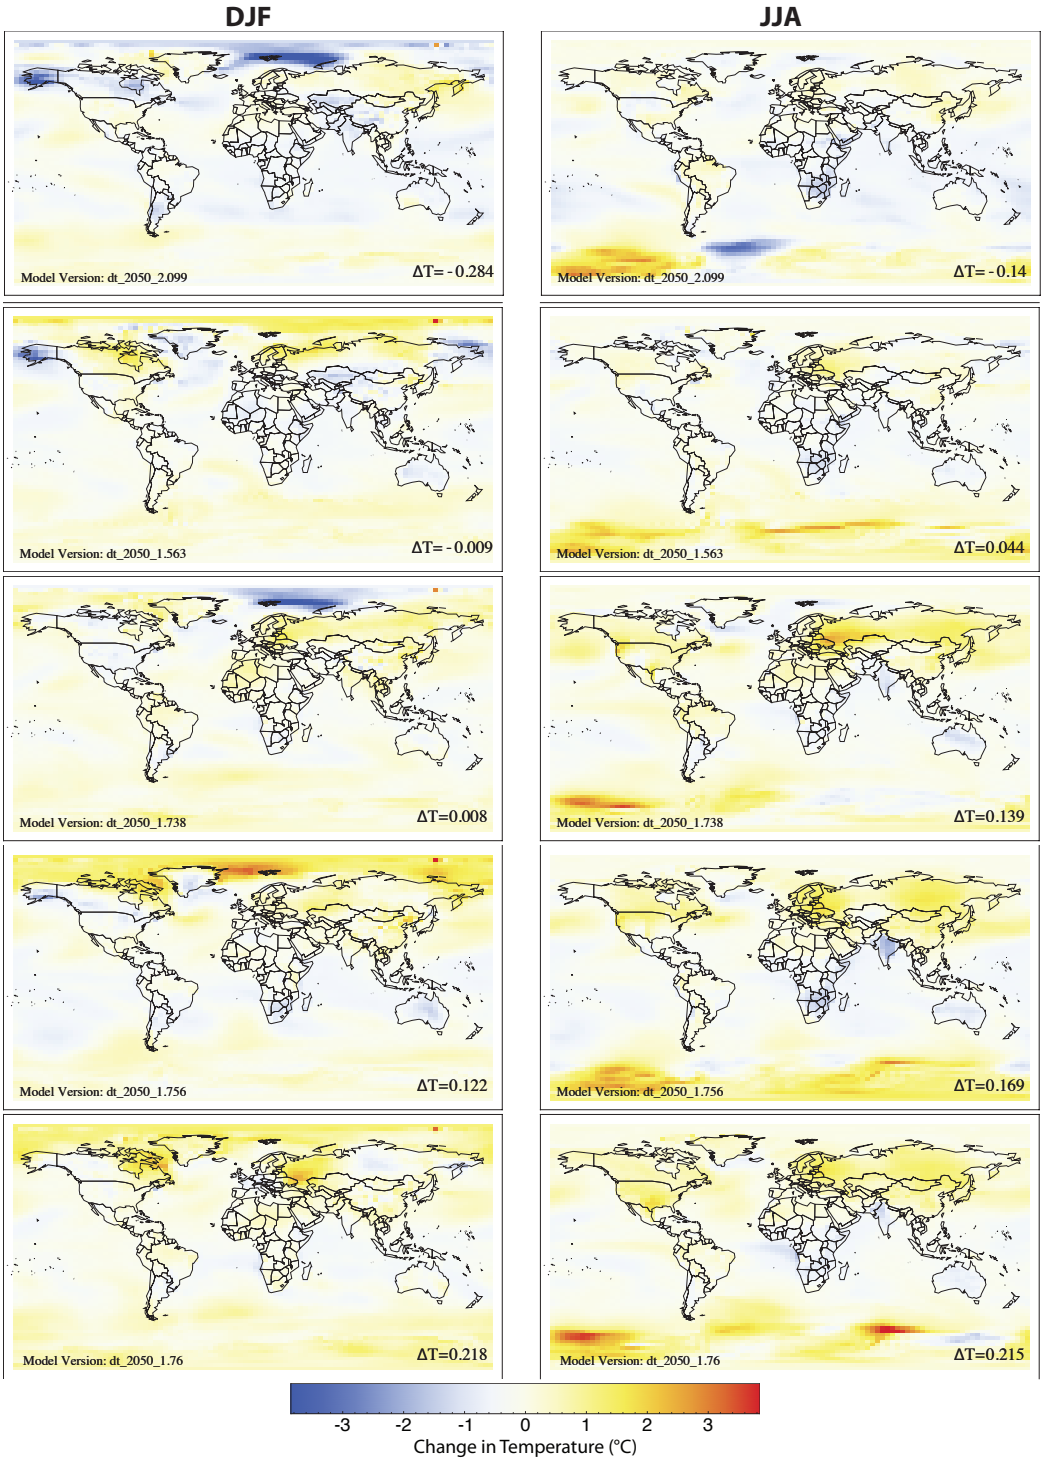

Figure 4.14: 20-Year Seasonal-Mean Near-Surface (1.5 m) Air Temperature Anomaly in °C (2050-1990) for high-SRM scenario. Maps for simulations using the standard physics model (top), followed by the four lowest temperature response models. Mean of three-member initial condition ensembles.

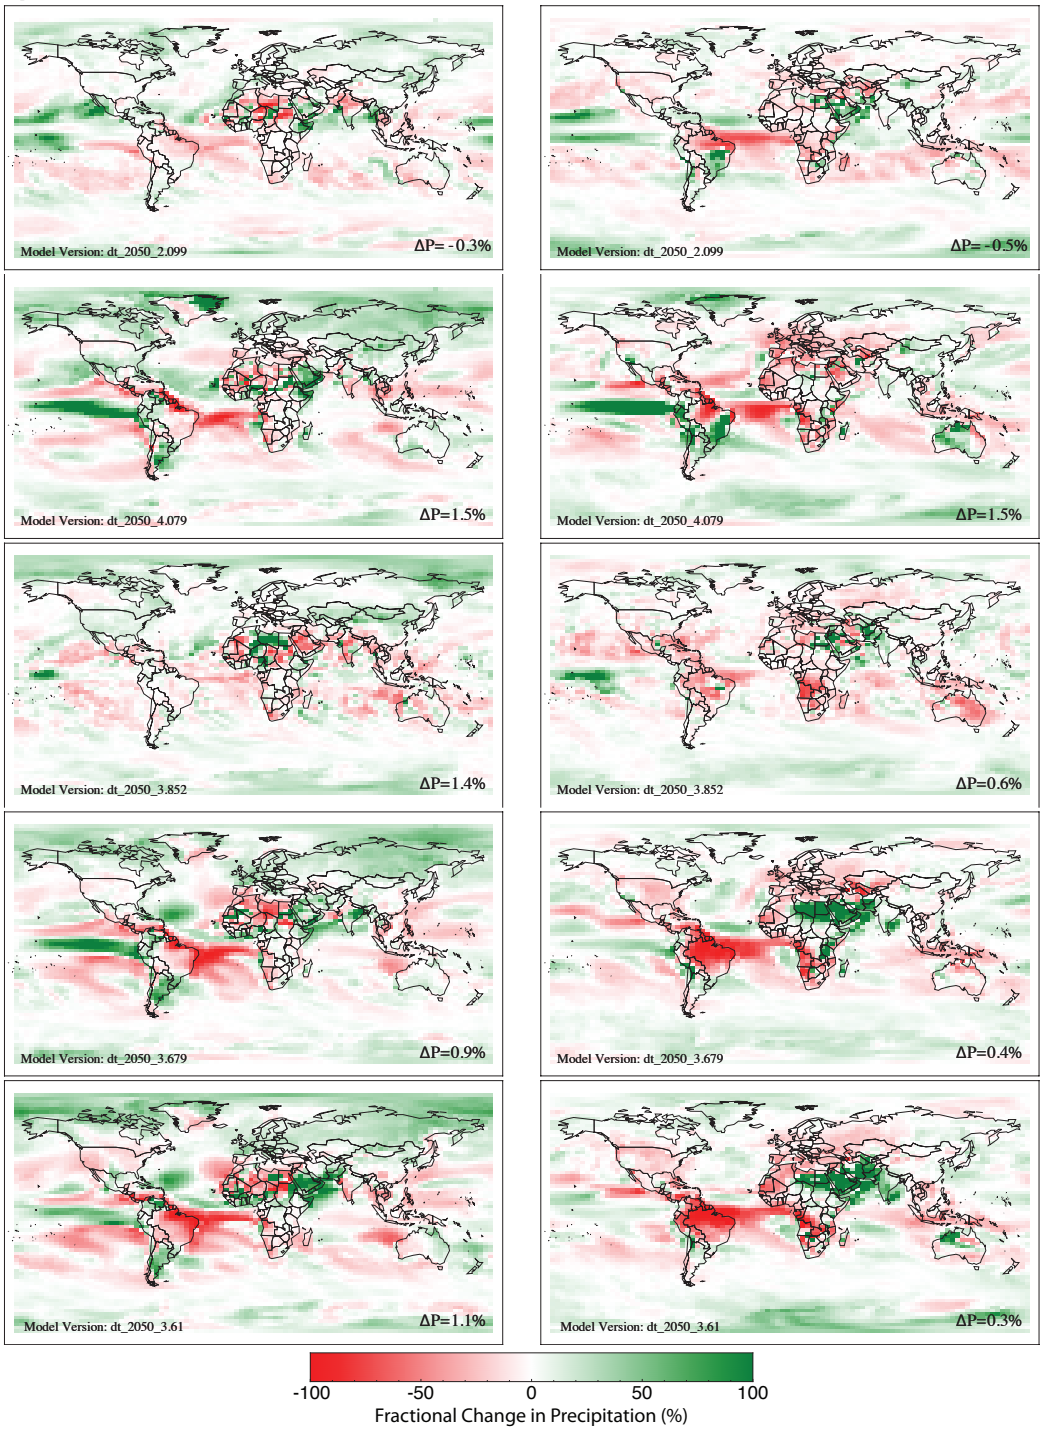

Figure 4.15: 20-Year Seasonal-Mean Percent Precipitation change (2050-1990) for medium-SRM scenario. Maps for simulations using the standard physics model (top), followed by the four highest temperature response models. Mean of three-member initial condition ensembles.

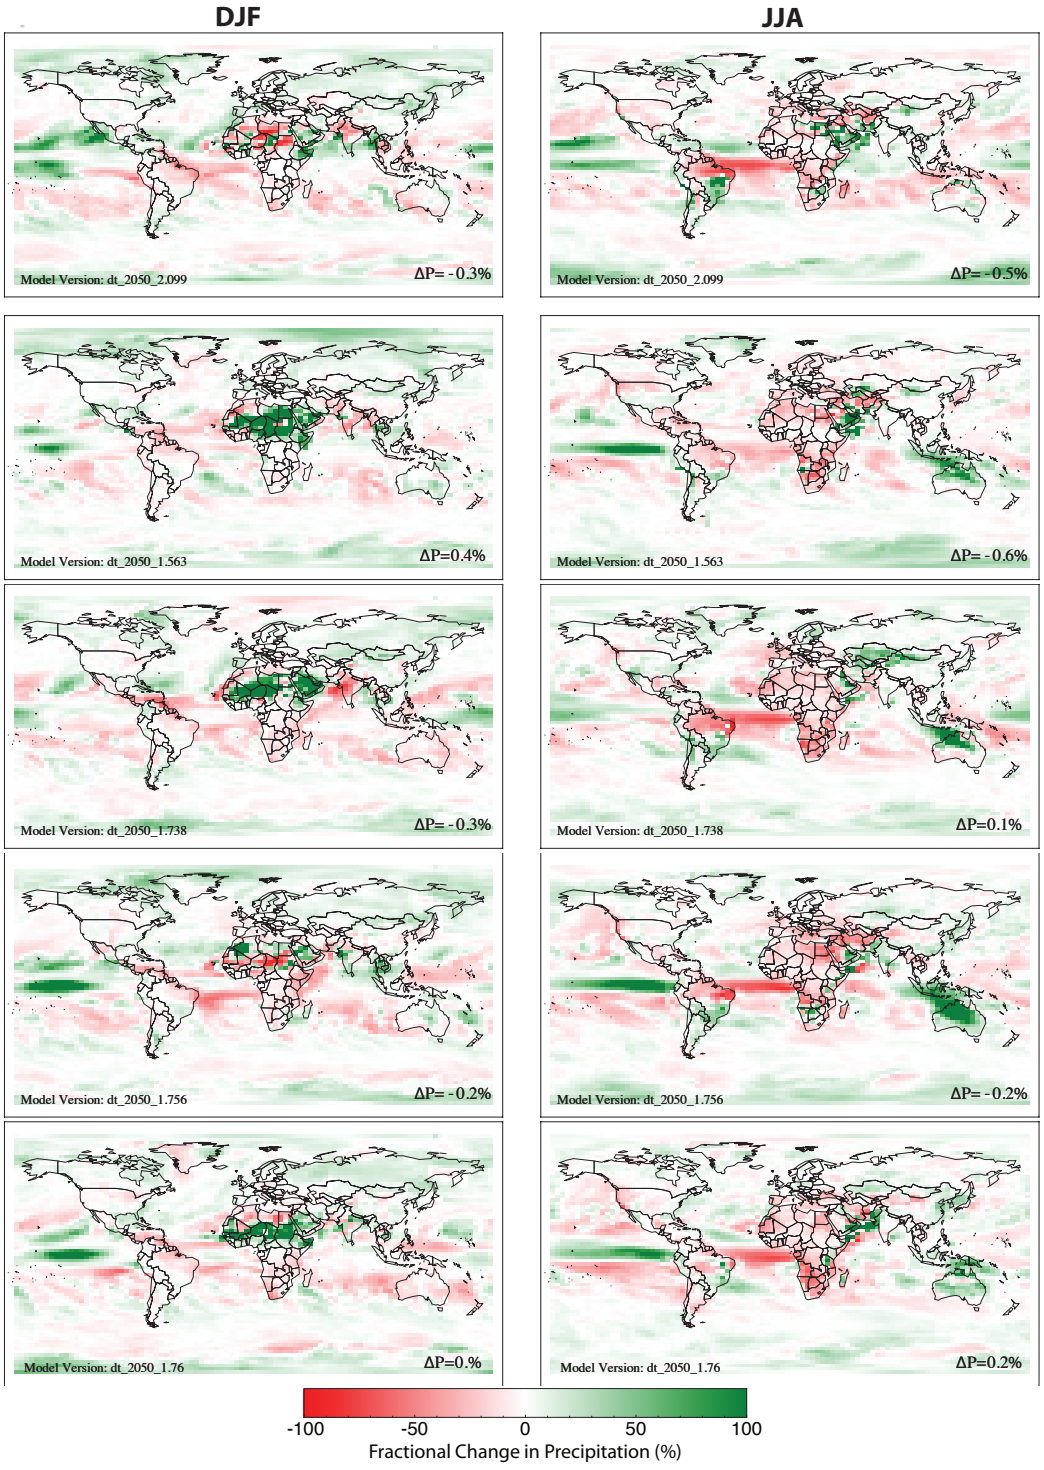

Figure 4.16: 20-Year Seasonal-Mean Percent Precipitation change (2050-1990) for medium-SRM scenario. Maps for simulations using the standard physics model (top), followed by the four lowest temperature response models. Mean of three-member initial condition ensembles.

Chapter 5

International Relations & Global Governance of SRM

There are several framings within which global governance issues surrounding SRM should be considered: as a part of a portfolio of strategies to address climate change, as a response to a global climate emergency, or as the tool of a rogue government worried only about their own regional climate problems.

As a part of a portfolio of strategies, together with emissions abatement and adaptation, SRM could be used to achieve an "optimal" response to climate change, thus minimizing net social cost or maximizing some other social objective. (Bickel and Lane, 2009; Moreno-Cruz and Keith, under review; Shepherd and Rayner, submitted) This framing is controversial given uncertainty about climate and ecological science (Matthews and Caldeira, 2007), the ethical issues involved in intentionally "engineering the planet," (Gardiner, 2011; Jamieson, 2006; Keith, 2000b) and the fact that there is no single global decision maker for whom an optimum can be defined (Morgan et al., 1999).

A somewhat less controversial argument is that we must study SRM now in order to be prepared, because the actual value of climate sensitivity could be considerably higher

than the current best estimates. Under this framing SRM would be deployed only as a "last resort" in order to avoid climate catastrophe. (Crutzen, 2006; Gardiner, 2011; Morgan and Ricke, 2010; Victor et al., 2009). [1]

A third framing to consider is that of unilateral deployment of SRM. Current international legal standards do not explicitly restrict any nation from engaging in stratospheric SRM. It is unclear that even if they did such restrictions could effectively block action, just as the Nuclear Non-Proliferation Treaty did not prevent non-signatories, India, Israel, North Korea and Pakistan, from developing nuclear weapons. Thus, technically, a single nation or even a very wealthy individual could take matters into their own hands. Such a scenario is in some ways related to the "last resort" multilateral approach above, but implications are different. With a multilateral or globally unified approach, the assumption is that the climate catastrophe to be averted would be sufficiently damaging to much of the world so that cooperation on the deployment of SRM could be beneficial to almost everyone. On the other hand, unilateral deployment could lead to the alleviation of climate change impacts for one group while imposing a mix of externalities on another without its consent or compensation. This is one of the reasons it is important to understand the potential harm and benefits associated with these activities not only globally, but in terms of relative regional incentives.

SRM presents two global governance problems that are related, but vastly different in character and scale. First, SRM is an emerging technology with a growing body of research on its potential implementation and effects, but still many associated risks and uncertainties. SRM research is, at this point, largely uncoordinated and unregulated. Norms and institutions are needed to ensure that SRM research — especially any field testing— is safe, transparent and in support of an equitable approach to solving the global climate change problem. Second, SRM is a potential game changer in terms of climate change

[1] These two paragraphs on framings were adapted from Ricke et al, 2011 (in draft)

policy negotiations. There is every indication from research thus far that this technology could stabilize global temperatures for a cost that is much less than mitigation (McClellan et al., 2010), making it an attractive but highly imperfect substitute or supplement for emissions reductions. In addition, its effectiveness in reducing the effects of rising greenhouse gases varies by region and is uncertain. These unique attributes of SRM could shift various actors' preferences and strategic approaches toward climate policies.

In this chapter, I aim to contextualize some of the science presented in the previous two chapters on regionally diverse climate responses to SRM forcings. Using basic international relations theory. I will characterize the issues raised by SRM, both as a risky emerging technology and a new problem to account for in the global climate change policy and negotiations process. Next, I discuss how SRM differs from other approaches to climate change in relation to uncertainty, which is paramount in the negotiations on global environmental policy, and analyze how SRM complicates already complex questions of whether and how to establish equitable climate change policies. Finally, I discuss the role of institutions and regimes through which foreign policies on SRM activities may be researched, tested or implemented.

5.1 International Relations Theories as Applied to SRM

There are three prevalent models of international relations in the early 21st century: *realism* pictures international relations as a series of strategic interactions by self-interested states; *liberalism* – aka *institutionalism*– posits that as democracy spreads, the world will become more peaceful and global policies more effective through the optimal design of multilateral institutions for cooperation; and *constructivism* presents changing relations and world order primarily as a products of shifting collective values. (Snyder, 2004) These different theoretical approaches to international relations theories imply different things

about the conditions under which SRM will be deployed, the likely sticking points in international negotiations on SRM as a part of global climate policy and what kind of global governance structure for SRM is plausible. The explanatory power of any one theory to explain the international relations of global climate change is limited, and thus, each model provides different insights into the dynamics of SRM governance problems.

(Neo)realism: Relative Gains Dominate Choices

Classical realism portrays the world as a collection of self-interested states which seek power (particularly military power) to leverage in relations with other states. More relevant to the discussion of global climate negotiations is the more recently developed neorealist theory, which emphasizes the distribution of state more complex "capabilities" in determining the outcome of negotiations. Central to neorealist models of international relations is the concept of hegemonic stability, in which the amount of international cooperation on an issue is determined by the extent to which a hegemon, or state with dominant power, dictates such cooperation. (Rowlands, 2001) Under the neorealist model, we expect that global climate policy will be set by superpowers, and in particular, one hegemonic superpower. Under realist theories, cooperation between states rarely results in optimal choices. Instead states make satisficing choices, ones that are "good enough" and strategically advantageous.

As the country with the most powerful military, largest GDP and responsible for the largest share of accumulated anthropogenic CO_2 in the atmosphere, the United States is the only nation today that could reasonably be considered a "climate hegemon," although China could achieve that status in the future. To a great extent global climate policy has been dominated by the U.S.'s inaction, including its refusal to ratify the Kyoto Protocol. In addition, U.S. concerns over the making large cuts while increasingly powerful devel-

oping countries such as India and China are held to less stringent standards has been a sticking point, suggesting that avoiding a competitive disadvantage is a major concern for the U.S. However, the considerable influence that other nations or groups have had in shaping the global debate on climate change – in particular the leadership role the European Union has assumed in regulation of greenhouse gas emissions– suggests that unless the U.S. were to become much more assertive, this model is limited in its application to this problem.

Assuming the neorealist model has some value in predicting climate change negotiations, several considerations on the role of SRM in these relations arise. To the extent that the U.S. is able to act as a climate policy hegemon, SRM is unlikely to be implemented if it is not in the interest of the United States. In addition, reaching international agreement on the specifics of SRM implementation will be complicated if powerful states have any significant divergent interests.

Institutionalism: Absolute Gains Matter More

Liberalism, also known as institutionalism, presents an optimistic view of the world as transitioning towards all democratic societies. With democracy, institutionalists posit, cooperation will increase and international policies will become increasingly optimal for all. Under the institutional model, action on climate change must happen when emitting nations all conclude it is in there best interest to do something rather than nothing. The fact that most countries have ratified the United Nations Framework Convention on Climate Change (UNFCCC) and participate in associated negotiations is proof that institutionalism has some relevance in international climate relations. (Rowlands, 2001) In contrast to the neorealist model, institutionalism would suggest SRM will be implemented if the global community as a whole will benefit from it and there are mechanisms (institu-

tions) to ensure effective negotiations.

Institutionalism suggests that strategic gains are less important than absolute ones. This dichotomy between relative and absolute gains could lead realists and institutionalists to draw very different conclusions about the prospects for cooperation based on the same data. For example, Figure 5.1 highlights the contrasts between realist and institutionalist worldviews as applied to Figure 3.9 in Chapter 3. Realism would suggest that divergent relative gains in determining the amount of SRM would make consensus difficult as China and India migrate out of their baseline climate circles and some regions are certain to become "losers" relative to others, whereas institutionalism would suggest cooperation is likely to continue considering the similarity of the two regions interests in comparison to a no-SRM alternative. In reality, both relative and absolute gains will likely be important in determining whether SRM is implemented.

Constructivism: Shared Values Matter Most

Constructivism presents a critique of both realism and institutionalism by de-emphasizing the role of state interests and asserting instead that the collective values of the agents of states and organizations are what motivate relations and institutions. (Viotti and Kauppi, 2010) Change is driven by "intellectual entrepreneurs who proselytize new ideas" (Snyder, 2004) and institutions arise as actors come together to codify shared values. One example of this is the emergence of salient international norms and institutions to protect basic human rights. Social upheavals around the world during and in the century after the Industrial revolution led to Abolition, legal changes to implement universal suffrage and workers' rights. Many of the norms that emerged from these social movements were codified in the UN Universal Declaration of Human Rights in 1948 and International Covenant on Civil and Political Rights in 1966. (Ishay, 2004) While such norms are still far from per-

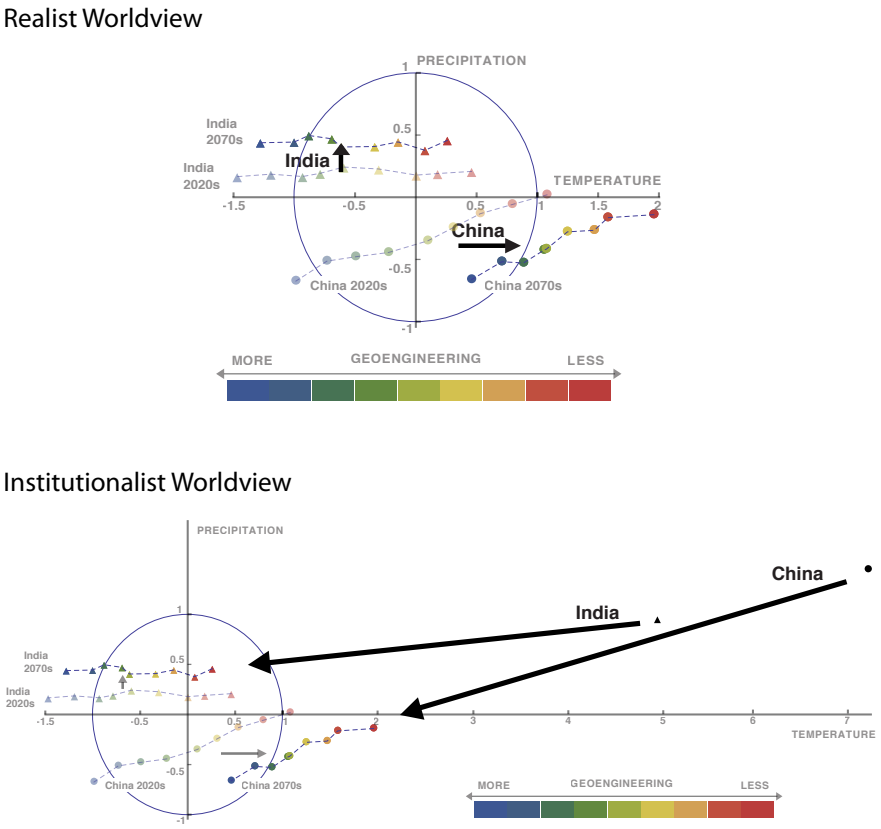

Figure 5.1: Realist versus Institutionalist Views of SRM, as applied to Figure 3.9.

fectly implemented in the international political arena today, the values the represent are widely espoused and provide justification for foreign policies of many prominent nations.

A central concept of modern constructivist relations theory is the concept of "name and shame" — the idea that actors, be they nations or nongovernmental organizations can change behavior of other actors by calling them out on action contrary to (Snyder, 2004). In order for such a mechanism for action to be effective, though, the values being violated must be pervasive enough to be true norms. One might hypothesize that under the constructivist model, meaningful action on climate change is unlikely to occur until a critical mass of international actors ascribe to a "sustainability norm" under which per-

sonal benefit from environmental externalities at the expense of people in other areas of the world or future generations is considered unacceptable.

In a world in which greenhouse gas (GHG) emissions reductions are predicated on the emergence of such norms, the existence of SRM could be counterproductive, as norms-based international interventions often have thresholds for action (Wheeler, 2000). Building upon existing norms of global justice, for example, suppose that human suffering linked to climate change must pass some threshold for humanitarian intervention in order for a a salient sustainability norm to emerge and provoke decisive political action to mitigate. Under such circumstances, it is possible that SRM designed to reduce suffering in the short term could ultimately increase net suffering.

Figures 5.2 -5.5 illustrate a thought exercise on how this could happen. Thresholds for humanitarian intervention are not determined based on ethical considerations alone, but also on practical considerations of feasibility (Wheeler, 2000). As such, this thought exercise is based on the assumption that the threshold for deploying SRM is lower than that for reducing GHG emissions, not for any normative reason, but because SRM deployment is: 1) cheaper than greenhouse gas reductions and 2) requires less international coordination (from a technical perspective). Figure 5.2 shows suffering from climate change over time as illustrated in a world where SRM is off the table. As climate change impacts increase, so too does human suffering. At the threshold point, suffering caused by climate change becomes so shameful to the global community that a sustainability norm becomes salient and political action to reduce GHG emissions occurs. Because of carbon inertia, even after such a norm becomes salient, suffering would likely continue to rise until CO_2concentrations and rate of climate change can be stabilized.

Figure 5.3 shows suffering from climate change over time as illustrated in a world where SRM is a possible means of reducing suffering at a lower cost than GHG reductions, or more quickly, meaning that the less imposing tradeoffs allow SRM to be implemented

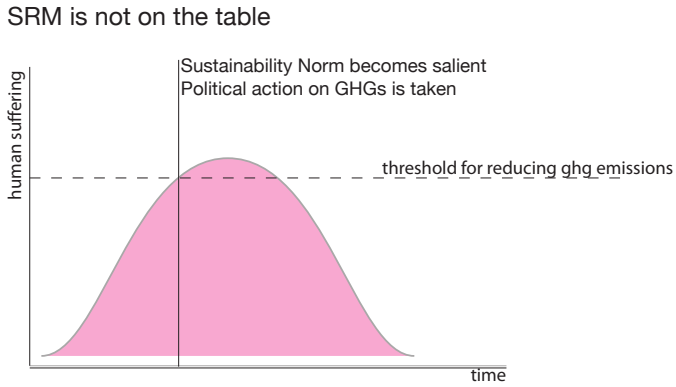

Figure 5.2: Constructivist thought exercise showing human suffering (or some other shameful impact of climate change) over time with and without SRM. Part 1: SRM is not on the table.

before a Sustainability Norm becomes salient. SRM slows the rate of suffering increases (or could perhaps even reverse it for a time)[2], but because it is an imperfect substitute for GHG reductions, eventually under such a scenario, total human suffering – as illustrated by the area under the curve – is actually much greater than it would have been in a world without SRM.[3] Note also that because CO_2concentrations would be higher by the time the threshold for reducing emissions is reached, it would likely take even longer to overcome the carbon inertia that prevents society from reducing associated suffering right away.

Of course, even with differential thresholds for intervention for SRM and GHG emissions reductions, SRM may not necessarily increase net suffering so dramatically as it appears to in Figure 5.3. For example, the human suffering threshold for reducing emissions could diminish over time due to the development of cheaper mitigation technologies or

[2]Unlike with GHG concentrations, the suffering reduction mechanisms of SRM could kick in immediately after the associated norm becomes salient because the timescales associated with his climate forcing are so much shorter.

[3]This argument, of course, assumes that on net, climate change will have a negative impact on peoples welfare. While climate change will almost certainly be unambiguously damaging to natural ecosystems, over the next century, its effects could be positive for some managed ecosystems, leading to increased food production, etc.

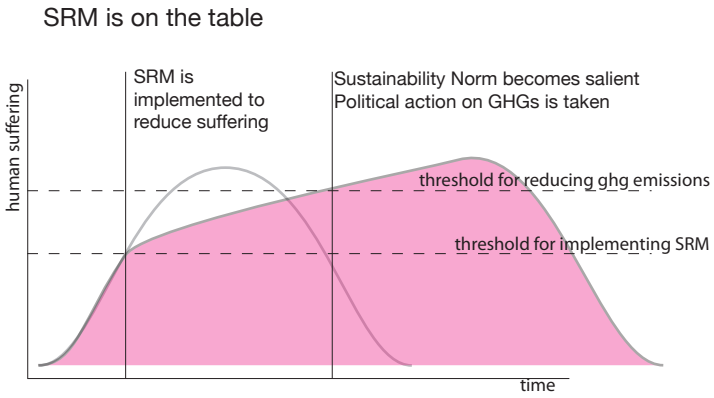

Figure 5.3: Constructivist thought exercise showing human suffering (or some other shameful impact of climate change) over time with and without SRM. Part 2: SRM is on the table.

changing social tolerances for suffering (Figure 5.4). In addition, the rate at which suffering could be reduced may be greater if the threshold for action is passed later (Figure 5.5). Even under such scenarios, though, when SRM is implemented, suffering is likely being pushed onto future generations. The constructivist lens reveals philosophical conundrums in the debate whether SRM would reduce the net negative impacts of climate change, even if it were determined that it could definitively reduce impacts at any given point in time.

Summary

Table 5.1 summarizes the three models of international relations and some of their implications for cooperation on climate change and deployment of SRM.

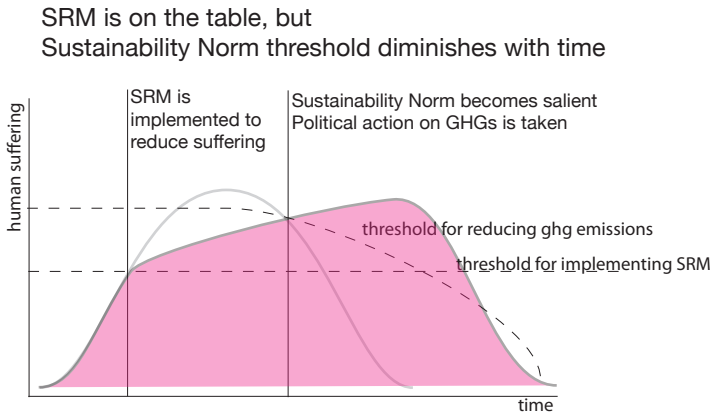

Figure 5.4: Constructivist thought exercise showing human suffering (or some other shameful impact of climate change) over time with and without SRM. Part 3: SRM is on the table, but the threshold for GHG reductions diminishes with time.

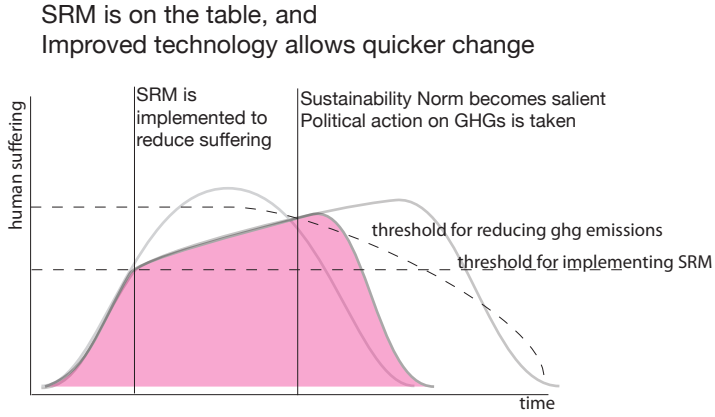

Figure 5.5: Constructivist thought exercise showing human suffering (or some other shameful impact of climate change) over time with and without SRM. Part 4: SRM is on the table, but the threshold for GHG reductions diminishes with time and capacity to quickly reduce effects of GHGs with abatement increases with time.

5.2 Role of Uncertainty

It is well known that scientific uncertainty significantly complicates issues of environmental decision making and global governance. (Viotti and Kauppi, 2010) One might

Table 5.1: Summary of International Relations Theories and Implications for Climate and SRM Policies.

	Neorealism	**Institutionalism**	**Constructivism**
Nutshell World View	Powerful states dominate decisions, action hinges on a strategic advantage of a hegemon	Democracy breeds cooperation. Good institutional design can lead to optimal choices	Shared Values lead to new norms; norms lead to reflective institutions
Cooperative Actions Reflect Decisions that...	Maximize the relative/strategic gains of the most powerful	Maximize net utility	Reflect the shared values of the decision makers (and, to varying extents, their societies)
Cooperation on climate change will happen when...	It benefits the United States or another hegemon/hegemonic coalition	The net benefits of cooperating exceed the costs and strong institutions are in place to coordinate action	A critical mass of actors in negotiations stare the same convictions about the actions required (e.g, a sustainability norm has emerged)
SRM will be deployed in situations where...	• the U.S., or a climate hegemon, is a winner in terms of relative SRM efficacy • it is strategically advantageous for the U.S. (as a dragger in negotiations, to buy more time/ to reduce suffering) to support deployment	• the absolute gains of deploying exceed the gains associated with alternative actions • institutions for compensating "losers" exist	Values of decision makers align such that alleviating human suffering/ protecting the environment or vulnerable resources/ etc quickly is more important than avoiding "tinkering with the planet"

imagine that international negotiations to implement SRM would be just a fraught as those to reduce emissions, but this may not necessarily be the case. Gardiner (2011) suggests this in his discussion of the problem of political inertia in the efforts to reduce global emissions and asks, regarding SRM: "if... we already have adequate scientific and tech-

nological solutions [to deal with climate change using mitigation], why assume that research on alternative solutions will help?" Part of the answer to this question would undoubtedly be that SRM offers the advantage of short timescales for a climate response. But additionally, if SRM reduces uncertainty in physical outcomes of climate change, achieving political consensus on its implementation may be easier than that for GHG emissions reductions.

Take, for example, Figure 5.6 that shows histograms of the global temperature and precipitation responses to GHG forcings with and without SRM for the perturbed physics ensemble analyzed in Chapter 4. The range of potential changes in temperature and precipitation is reduced with SRM, a response observed at the regional level as well. By reducing the temperature component of climate perturbation, uncertainty in the range of response of all climate indicators is reduced.

The role of uncertainty in establishing global political consensus has often been illustrated by comparing the uncertainties associated with ozone-depleting chlorofluorocarbon (CFC) emissions and those of climate-changing GHG emissions. The Vienna Convention and Montreal Protocol are generally hailed as the most successful major global environmental regime. (Sprinz, 2001) Efforts to establish effective global policy on protecting the ozone layer benefited from rising certainty in the beliefs of the effects of reducing CFC emissions and highly certain preferences for outcomes among negotiating parties. On the other hand, the regional effects and preferences associated with rising atmospheric GHGs are both highly uncertain and both of these uncertainties reduce the prospects of global political consensus. Figure 5.7 adapts a J.D. Thompson-style decision matrix to illustrate this point and how SRM might shift the climate change point. In this plot, political consensus is easiest in the upper left-hand quadrant and most difficult in the lower right. (Viotti and Kauppi, 2010, pg. 137)

The type of effect observed in figure 5.6 suggests that SRM could shift the climate

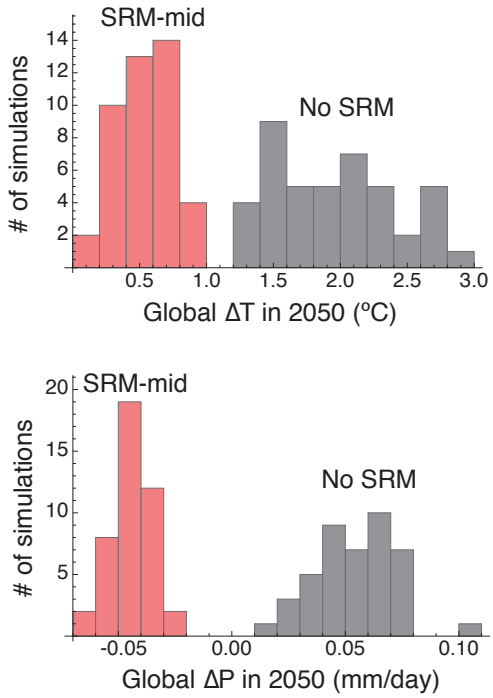

Figure 5.6: Range of outcomes with and without SRM from the Experiment 2 PPE (see previous chapter). The top figure shows a histogram of temperature outcomes in the ensemble with SRM (in red) and without (black). The bottom figure shows the same for precipitation. Note that in both cases the spread of predicted values narrows with SRM.

change problem upwards in the y-dimension of Figure 5.7, toward greater certainty in beliefs about environmental cause-effect relations. In addition, the mechanisms by which SRM changes the climate are simpler and have been observed (by way of the volcanic analog). The distinction between true scientific certainty about cause and effect and "certainty in beliefs" is an important one. Victor (2011) points out that scientific uncertainty surrounding the effects of CFCs on the ozone layer was actually quite high at the time of the formulation of the Montreal Protocol and that interests, in particular from the private sector, played a large role as well. However, the observations of the Antarctic Ozone Hole two years prior still spurred action from decision makers and industry, even if scientists

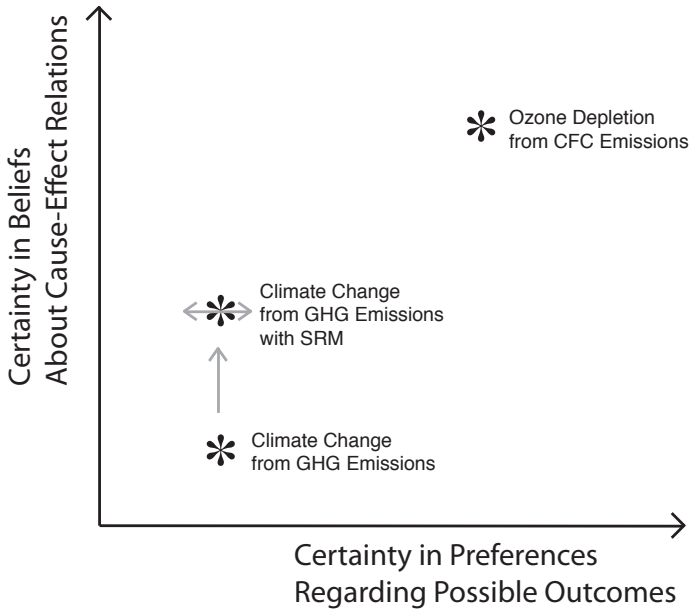

Figure 5.7: J.D. Thompson-style decision-making matrix showing dimensions of uncertainty as applied to international environmental political decisions. The effect of SRM on the position of the climate change problem is proposed as upward on the cause-effect access and unknown on the preferences axis.

had not yet resolved the mechanisms which created it.

On the other hand, the effect of SRM on the points position in preference space is hard to ascertain. Given the diversity of regional economic goals, capabilities and satisfaction with current climate states, SRM may increase uncertainty in this dimension. In addition, SRM introduces and element of choice (i.e., an ability to "set the thermostat").

5.3 The Need for Research

The scientific uncertainty about SRM's impacts and the lack of international legal frameworks to address it yield an unsettling policy void. A broad and solid foundation of research will help on two fronts. First, it will transform the discussion about geoengineer-

ing from an abstract debate into one focused on real risk assessment. And second, better understanding of the dangers will help in crafting the norms that should govern testing of geoengineering systems, provide information in the case of an attempted unilateral deployment of SRM and guide collective deployment in the event of a sudden global climate disaster.

There are a number of reasons why investing in SRM research could be harmful, including:

- the risk of "lock-in". Research initiatives can build up institutions that create momentum for the deployment of a technology regardless of the results of the research itself;
- the potential for moral hazard. If SRM can help insure against the worst impacts of climate change, we may take action to mitigate it more slowly that we would otherwise; and
- the diversion of resources from other climate science and research and development of mitigation technologies.

Despite these concerns, the risks of not investing in SRM research are even greater. The remainder of the section, adapted from Morgan and Ricke (2010), attempts to systematically make the case for this assertion.

Figure 5.8 provides a simplified illustration of the decisions about whether and what research to conduct and what might be learned from that research. Along the left side of the graphic is a time line that illustrates that up until now very little research of any kind has been conducted on this topic. If we decide today to engage in a program of computer and lab studies– and some such programs have already begum – we might then, on the basis of what is learned in those studies, choose to conduct a set of field studies. The

Figure simplifies what might be learned from those studies into three broad outcomes. At worst, we could learn that SRM would not work to reduce the negative effects of global warming or would result in unmitigated climate or ecological disasters. At best, research may indicate that SRM can be done easily and inexpensively, both in terms of direct costs and externalities. Alternatively, we may learn that SRM is more expensive to implement than anticipated or would have previously unforeseen negative side effects. In either case – or in any case in between – such information would allow the world to more realistically compare costs and side effects of an SRM-modified world with those associated with one with no SRM.

Suppose now that at some time in the future a country or region finds itself facing a serious local or regional problem caused by climate change. While the rest of the world might take action to aid that region (for example, by providing food aid in the event of a profound drought), suppose that instead the nation or region chooses to take matters into its own hands and unilaterally "solve" its problem by engaging in SRM. In this case, it would impose any associated effects and externalities upon the rest of the world. The "SCENARIO ONE" row of Figure 5.8 summarizes in very simple terms the situation in which the rest of the world would find itself as a function of what had been learned as a result of the previous program of research. In the case of an anticipated "unmitigated disaster", the world should have taken action to formally restrict such activities. If the research program has characterized the costs and risks associated with SRM, the international community will be in a position to take an informed stand in opposition to such unilateral action, on the grounds that, in the view of the authors, no single nation or region should have the right to unilaterally impose the externalities on the rest of the world. If no research is conducted, any opposition would be uninformed and therefore less legitimate.

Finally, the "SCENARIO TWO" line illustrates the situation in which, in the future,

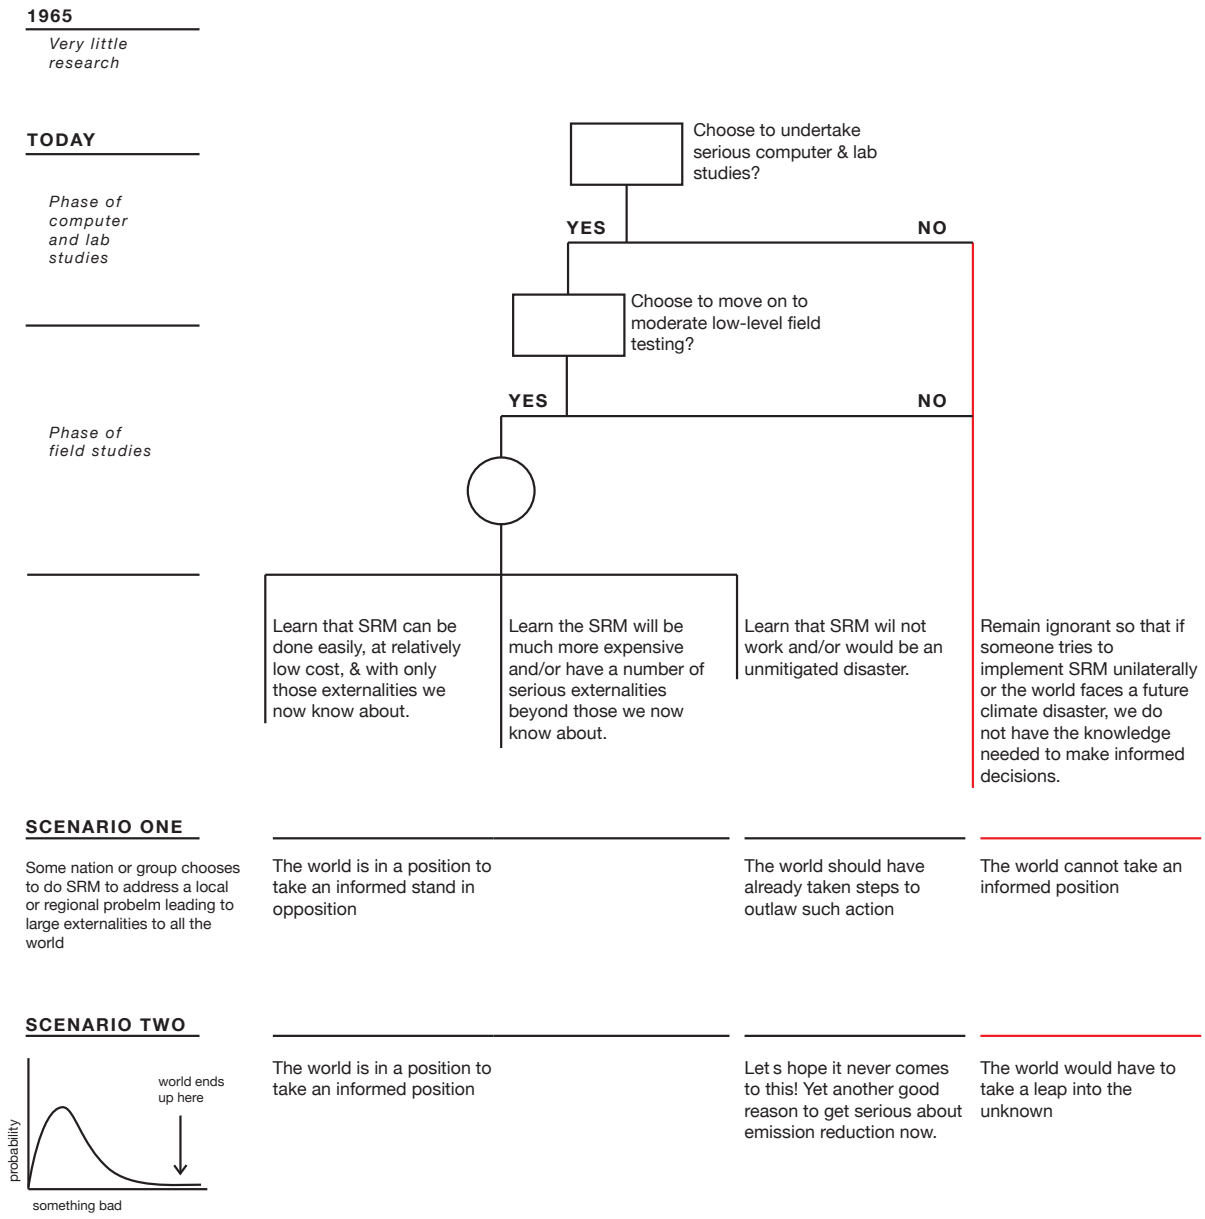

Figure 5.8: SRM Research Timeline and Decision Tree,adapted from Morgan and Ricke (2010).

the world finds itself facing a very serious global climate disaster. In this case, if the research has demonstrated that SRM is feasible, the world would be in a position to weigh

the relative costs of engaging in SRM or suffering the consequences of the climate change.

The probability of this extreme outcome may not be entirely independent of whether or not the world does research on SRM. As suggested above, one reason that many scientists have been reluctant to engage in SRM research is concern that knowing more about it may increase the chance that someone relies on it as an alternative to reducing emissions of CO_2 and other greenhouse gases. However, there is also the possibility that if we conduct research that shows that SRM will not work, and/or would be an unmitigated disaster, which might induce a more serious global effort to engage in dramatic reduction of emissions. And, without research, if a climate disaster does arise, the world may end up faced with the decision to take a leap into the unknown and deploy SRM despite ignorance of the consequences.

5.4 Equity Challenges

Equity considerations are important in the development of effective and legitimate global governance of both SRM research and deployment. In the case of research, the challenge is to build a body of knowledge which allows the assessment of the impacts and risks to all nations, despite the concentration of scientific resources in the wealthier ones. In the case of deployment, the questions of who gets to decide when, how or to what extent SRM is carried out are fundamental to the debate about whether using this technology to counteract global warming is politically feasible, wise or ethical.

The addition of SRM to the climate policy mix adds a layer of complexity to the diversity of interests climate change creates. The inequitable distribution of vulnerabilities and abatement costs already results in a negotiating environment of diverse international (and intra-national, and transnational) interests. Figure 5.2 from Sprinz and Vaahtoranta (1994) highlights some of the broad interest groups that emerge from this

heterogeneous distribution of GHG emissions costs and benefits. This framing suggests that those with low abatement costs and high ecological vulnerability will be leaders in the push for strong global agreements on abatement, whereas those with high abatement costs and low ecological vulnerability will try to delay such agreements. With SRM, a new dimension needs to be considered contrasting interest groups for whom SRM is highly effective and those for whom SRM is less effective. This dimension could create a new significant group of "Intermediates" for whom SRM is effective who may become "SRM Pushers" or "Intermediates" for whom it is not effective who become more like "Abatement Pushers."

Table 5.2: Interest Matrix for Climate Change Policy Considerations. (Sprinz and Vaahtoranta, 1994)

	High Ecological Vulnerability	Low Ecological Vulnerability
Low Cost of Abatement	Pushers	Bystanders
High Cost of Abatement	Intermediates	Draggers

Add to this diversity of interests traditional ethical considerations of differential responsibility for accumulated GHG emissions, differential economic capabilities, developmental statuses and intergenerational effects, and designing institutions or regimes to promote equity becomes a complex task indeed. In addition, "universalism" in international environmental policymaking, the attempt to include all players for fairness sake, can produce institutions that are ineffective, and not necessarily more legitimate. (Victor, 2011)

Wiegandt (2001) characterizes three major equity challenges in international negotiations on climate change. First, there is the challenge of determining what is "fair" in terms of the distributive outcomes of a given action (or inaction). Second, parties need to agree on rules to achieve such a fair outcome. Third, there needs to be an effective system for enforcing these rules and monitoring what they achieve. Addressing these challenges

in the case of research and deployment of SRM requires different approaches.

In both cases, the challenge of determining what constitutes fair is a debate that requires input from a broad set of stakeholders, in particular a set that represents the interests of both developing and industrialized nations, as each stakeholder may have significantly different environmental priorities that influence that definition. In the case of governance of research, the second and third challenges which will require significant technical focuses, inclusiveness may not be the best pathway to ensuring research programs that efficiently generate knowledge to support equitable outcomes.

Considering the disparate distribution of research capabilities between developed and developing countries it would be disingenuous to promote equity in SRM research governance by simply giving developing nations a seat at the institutional table. Rather, it is important to promote research programs that incentivizes the best research institutions to consider the effects of SRM on all regions and ecosystems. Likewise enforcing the rules for equitable research and monitoring its outcomes is likely to be most effective when the institution(s) policing have technical expertise first and foremost (with secondary oversight from the same type of stakeholders that determined what is fair in step 1.)

On the other hand, equity in SRM deployment may mean giving all stakeholders a seat at the table and the best possible information with which to negotiate their interests themselves while establishing rules. Unlike GHG emissions reductions, or SRM research, full scale implementation of SRM would constitute a deliberate modification of the environments of all nations. Considering this direct effect, rules and systems to govern implementation may be to be viewed as legitimate without greater inclusiveness. With such high stakes and (probably) ties to other climate change governance regimes, enforcement and monitoring will also necessarily require oversight from broader institutions.

5.5 Governance Considerations

International institutions designed to tackle the climate change problem, specifically the UNFCCC, have not thus far been very successful. Universal participation has made the UNFCCC process cumbersome and legally binding targets and time tables have been simultaneously too vague to implement and unfeasible to achieve. Morgan (2000), Victor (2011) and others have proposed new approaches to international climate policy that begin smaller groups of key players. These smaller groups would forge emissions reductions agreements that become more comprehensive and inclusive with time and proven success.

This limited players approach could apply well to the problem of governing SRM as well. If certain "great powers" could come together to agree not to unilaterally geoengineer among themselves, it would would go a long way towards constraining the behaviors of other types of potential unilateral geoengineers as well. It would also provide an institutional base from which negotiations for collective deployment could begin.

Beyond the concerns over equity, discussed and contrasted in the section above, the major challenges associated with governance of SRM research and SRM deployment are different and should not be conflated. Establishing salient norms for coordination, transparency, and developing research program structures that allow effective allocation of limited resources will be critical challenges in establishing research governance. Preventing unilateral deployment as well as ensuring explicit ties between deployment and GHG emissions reductions will prove to be the biggest challenges facing the development of effective deployment governance.

Defining the Difference between Research and Deployment

In the international relations of climate change mitigation and adaptation, evidence suggests that while the epistemic community has had a strong influence on agenda setting — i.e., calling attention to and educating decision makers about the problem — they have had almost no influence on the political process of negotiating and creating regimes to address it. (Rowlands, 2001) While SRM research has implications for the entire global community, the institutions that govern it will effect primarily members of the epistemic community. It stands to reason that the most effective norms for research would be those that reflect the values of scientists and other researchers. This is not to suggest that the institutions for governance of technical research should not reflect the values of societies funding them, but only that they should prioritize the systematic reduction of the major uncertainties surrounding the risks and impacts associated with these technologies over the consideration of their geopolitical implications.

However, if institutions for research and deployment are separate, there must be a clear definition of where research ends and deployment begins. In Morgan and Ricke (2010), we proposed an "allowable zone" approach to setting *de minimis* criteria under which field research activities can proceed with minimal international oversight (in addition to observing the types of norms of discussed in the next section). An allowable zone may be defined by things like the experiment's proposed magnitude and duration of radiative forcing and the total (estimated) ozone destruction associated with it.

Establishing reasonable *de minimis* standards should be a focus of early computer and lab-based SRM research programs. It may not be very difficult. "Process" type field studies that would evaluate aerosol processes, effects of atmospheric chemistry and the efficacy of aerosol dispersal technologies would likely require orders of magnitude less materials and radiative forcing than "climate response" type experiments that evaluate the global

and regional climate effects of SRM. Since the majority of the types of field research that may be done in the near future fall into the process category, anyway, it may make sense to just lump any hypothetical climate response field testing in with deployment which, arguably, it is anyway. (Keith et al., 2010; Robock et al., 2010).

Governing Research

Several efforts to establish preliminary guidelines for SRM research have been made by organizations primarily composed of elite scientists and other academics. For example, the Asilomar International Conference on Climate Intervention Technologies in 2010, the largest meeting convened with the aim of developing SRM research norms to date, produced five suggested principles (ASOC, 2010):

- SRM research should focus on the "collective well-being of humankind and the environment";
- Mechanisms be established for oversight that specify responsibility and liability for research activities;
- Research plans should be transparent and results should be freely available;
- Research is contingent upon iterative and independent assessments of its progress; and
- Opportunities are provided for public oversight and consent.

These principles are closely related to the "Oxford Principles" generated by a group of UK scholars in response to a House of Commons inquiry in 2009. (Rayner et al., 2009) There is a general consensus among those thinking about research governance today that crafting non-binding norms and then inculcating the research community is the most

sensible course of action at this stage. This makes sense because major restrictions on research activities could just result in no research being conducted at all, or research that takes place by way of loopholes.

Some of the norms that have been suggested aren't particularly substantive. Who decides what collective wellbeing is? Most scientists probably believe that their research promotes such collective well-being, but their ideas about what this is are obviously shaped by the society they live in (usually a rich industrial one). Similarly, it is worth asking *which public* is realistically going to be involved in oversight or providing consent for research activities? Again, this will probably be the public of the rich, industrial societies (that are arguably best equipped to adapt to climate change and may not need SRM as badly.)

Other norms above may not be particularly relevant to what becomes regarded as allowable research. If research activities are constrained by *de miminis* standards, establishing responsibility and liability for climate-related effects is a problem better associated more with deployment than research. Liability for the effects of smaller scale field tests would be better handled by more local forms of governance (where existing opportunities for things like public oversight can be taken advantage of as well).

In contrast, establishing substantive norms on an allowable research zone and creating systems for transparency, cooperation and iterative and independent research assessments may require novel institutions. IPCC may be able to help with asessment and synthesis, but is not appropriate for any of the other tasks. An organization like the International Council for Science (ICSU) that represents and brings together elite scientific organizations from around the world would be a good choice for coordinating the process. Victor et al. (2009).

Negotiating and Governing Deployment

Because it could be implemented cheaply, quickly and unilaterally Keohane and Victor (2010) characterize the institutional challenge associated with SRM as being the exact opposite to that of other climate policies, "how to make it *more difficult* rather than easier to act."

Some international relations theories imply that the prospects for successful unilateral deployment of SRM are quite small, as the incentives to bring in partners with similar interests to attempt to increase the legitimacy of the action would be large. However, under such a scenario, it is unclear how such a limited multilateral deployment would be substantively different than a unilateral one. For example, while the U.S. drew upon such partnerships in the lead up to its 2003 invasion of Iraq, most still consider that war a unilateral action.

Some have argued that unilateral SRM is technically infeasible because acting alone might result in simultaneous deployments that interact or counteract in disastrous ways, or because after unilateral deployment other nations might attempt to counteract SRM by intentionally injecting potent GHGs or black carbon to counter its effects. (Horton, 2011) These scenarios seem unlikely. Once someone starts deployment, global observational infrastructure will likely inform everyone rapidly. And just because SRM could be counteracted in principle with some dramatic counterengineering doesn't mean it is realistic. And if the nation or group of nations that would like to counteract a rogue geoengineer's actions is powerful enough, it's far more likely that they would halt SRM simply by finding and destroying the infrastructure being used for deployment (and waiting for it's temporary effects to dissipate) rather than imposing an additional global environmental harm. If the counterengineering is not powerful enough, the rogue geoengineer could just respond in kind.

Ideally, the unilateral deployment of SRM should be banned. Practically, how a ban could be effectively achieved is unclear. There seem to be several types of potential unilateral geoengineers: great powers, rogue states and desperate states. (While I use the word "states," technically they could be non-state actors.) Great powers, such as the United States, China or Russia, are states that have a large enough combination of economic, political and military capacity that their actions would be difficult to reign in even with coordinated international action. Rogue states, such as North Korea, who operate in enough isolation from the rest of the international community and possess some dangerous military capabilities could also be difficult to prevent from engaging in unilateral deployment, as sanctions and conventional diplomacy tools are ineffective with such nations. Desperate states would be nations or groups that prove particularly vulnerable to climate change earlier than the rest of the world. These are states that would be highly unlikely to pursue unilateral actions against other states under any other circumstances, but whom the impacts of climate changes have made desperate. For example, small Pacific islands in Polynesia or Micronesia, are considered particularly vulnerable to sea level rise and might be candidates in this category.

Institutionally, the best way to prevent unilateral deployment of SRM may be to address just the first type of states. If great powers could reach consensus and even codify their opposition to unilateral action, it would constrain the other two types of potential unilateral actors as well. It would put the world in a better position to take action against rogue state geoengineers (by force if necessary) and it would make desperate states less likely to seriously consider unilateral action at all. For example: many of the vulnerable nations that are candidate "desperates" simply wouldn't be able to afford taking a stand against allied great powers. The economies of Polynesia and Micronesia are heavily dependent on foreign aid, military installations and tourism from those powers. (cia, 2009)

Getting a constituency of major powers to agree not to unilaterally deploy SRM would

not necessarily be an easy task. But it would be much simpler than trying to prevent all potential types of unilateral SRM through a universal process, while simultaneously not stifling scientific research on SRM. And as a similar approach was taken to international governance of ozone-depleting substances through the Montreal Protocol, there is some precedent for success. (Benedick, 2011)

A great powers alliance to prevent unilateral geoengineering, in turn, could be a institutional base for negotiating collective deployment of SRM. If SRM is deployed it must be linked with GHG emissions reductions, even if it is used in the "emergency response" capacity. If not, SRM will be nothing but a tool for pushing the impacts (or the emergency) of climate change off onto future generations while simultaneously damaging the ocean ecosystems (via prolonged elevated CO_2 concentrations), the ozone layer, or any other number of negative impacts associated with SRM. Explicit ties to emissions reductions may be complicated to implement given the conflict between certain reasons SRM may be implemented (i.e., a climate emergency requiring immediate action) and the endless gridlock associated with mitigation negotiations. As such, guaranteeing this link may be the other most important institutional challenge to preemptively begin to address.

5.6 Closing Thoughts

This chapter has presented a discussion of some of the international relations and global governance issues surrounding stratospheric SRM. In the past four years, the amount of scientific research published on potential effects of SRM has exploded. At the same time, a discussion about the regimes that need to be developed to govern this research and potential future deployment of SRM technologies has gained momentum. There has some significant discussion about the ethics of climate engineering (Gardiner, 2010, 2011; Jamieson, 2006; Morrow et al., 2009), but as Gardiner (2011) points out, just the serious

consideration of SRM "suggests that society is already in the grip of a moral failure of spectacular scope and import" by failing to mitigate. What is most ethical is not often what happens. Developing a normative approaches to SRM governance based on ethical precedents, theories about optimal outcomes and other laudable goals is obviously important as a benchmark to understand what is possible. Reconciling ideas about ethics and optima with analyses of what is probable or feasible, given political reality, is necessary to give proposed norms' and institutions' any chance of success.

Chapter 6

Conclusions & Future Work

In this thesis, I used large ensembles of simulations of a general circulation model to explore the basic dynamics of regional inequities in a world where greenhouse gas (GHG) driven climate change is compensated for using SRM. Different regions respond to SRM forcings in different ways. The first experiment (Chapter 3) systematically demonstrated some intuitive, but novel, results: that given some criterion for preferences, regions prefer different amounts of SRM; and that, as SRM is increased with rising GHG concentrations, these regional preference diverge over time.

Results from a second experiment (Chapter 4), demonstrated that these basic conclusions appear to be robust to parametric uncertainty in the model. The second experiment also showed that the effectiveness of SRM by various measures varies with the climate sensitivity or forecast warming the the model, an important consideration depending on what circumstances you believe SRM implementation is contingent upon. The absolute efficacy of SRM in returning regional climates to their past states decreases in higher climate sensitivity worlds, but the ratios of these residual regional anomalies after SRM is applied to the no-SRM alternative also decrease. The magnitude of these differences between high and low climate sensitivities depends on the assumptions you make about

the damage functions associated with impacts, but the basic trends appear robust.

The conclusions about how appealing SRM may be in a higher climate sensitivity world are complicated. Viewed from 2011, SRM looks less effective in the situations in which it is most likely to be used (i.e. when GHG emissions are unrestrained and climate sensitivity is higher-than-expected). However, viewed from a future where climate changes have already taken place, SRM looks most effective in reducing risks and damages associated with the situation at hand.

There is much more work on SRM to be done before any conclusions on whether or not to deploy it are reached, but this research suggests that regional inequities in climate response are probably not the main impediment to its effective implementation. While SRM can never perfectly correct for regional climate change, it almost always reduces (rather than exacerbates) changes to regional temperature and precipitation and greatly reduces the rate of temperature change.

SRM does nothing to prevent further acidification of the oceans [1] and this fact, coupled with the fact that SRM can only slow, never halt, changes to regional climate states, makes SRM completely untenable as a long-term solution to the problems caused by rising GHGs in the atmosphere. Considering the slow progress society has made towards reducing emissions, however, it would be wrong to not consider the potential benefits SRM technologies may confer in reducing impacts to buy time for both mitigation and adaptation. There are many risks associated with SRM that require further consideration and study. The following section explores some of the research needs that are derivative of the work presented in this thesis.

[1] Because SRM reduces surface temperatures, it changes the ocean and biosphere's rates of uptake of CO_2. This, and lower SSTs, would have some impact on ocean acidification, but that effect would be small.

6.1 Future Work

Additional PPE Analyses

The "climate*prediction*.netGeoengineering Experiment" has produced a very large data set that merits further exploration. As a part of the analyses presented in this document, I have conducted only a cursory examination of the certain diagnostics. The focus of the data analyses I have conducted so far have primarily been on near-surface air temperature and precipitation rates, both at the global and regional level. Additional in-depth analyses of the PPE output should take advantage of other output data, including:

- Runoff and soil moisture: In my analysis of the standard physics experiment, subsurface runoff was presented as an alternative proxy (beyond precipitation) for the effect of SRM scenarios on regional hydrological cycles because it is a function of soil properties, temperature and precipitation, and therefore may better portray information about impacts. Further analyses of the PPE data could delve deeper into such impacts-relevant hydrological indicators.

- Land-Ocean Heat Fluxes: In the standard physics experiment simulations, the global-mean precipitation response to SRM forcings over time follows a simple trend: as greenhouse gases and SRM increase, and global temperatures remain stable, precipitation decreases. Regional trends over land, however, are often quite different. Several regions, as well as land-only global-mean precipitation, do not exhibit trends consistent with the global mean—while regional temperature remains constant or changes steadily over time, precipitation trends have a kink around model-year 2050. One hypothesis for this phenomenon is that heat flux from oceans to land is reduced. In order to tease out the causes of differences in regional precipitation response, changes to winds and heat fluxes along land-sea boundaries in the

gridded data output need to be examined.

- Other radiative and heat fluxes: As with land-ocean heat fluxes, looking at the distribution of changes to short- and longwave radiative fluxes and latent and sensible heat fluxes will allow us to better understand the mechanisms behind regional precipitation trends. A first step would be looking at global-mean precipitation changes compared to transient radiative and heat fluxes could test results from Allen and Ingram (2002), Yang et al. (2003) and Andrews et al. (2009) in the context of temperature stabilization with SRM; seeing how compensation of CO_2 forcings with shortwave perturbations affect latent heat fluxes, radiative cooling and ocean heat uptake in models with different climate sensitivities.

Experiment 2 has PPE output for additional SRM scenarios. Figure 6.1 shows the full set of SRM scenarios that were tested in the PPE experiment.

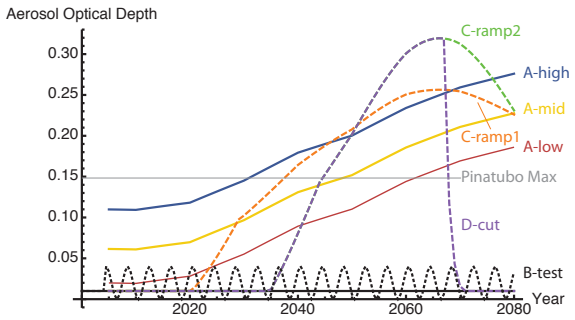

Figure 6.1: Several additional SRM scenarios were simulated as a part of the PPE experiment, including "B-test," "C-ramp1," "C-ramp2," and "D-cut." Data from the simulations using these scenarios have not been analyzed yet.

Testing the Climate Effects of SRM

One question yet to be fully addressed by researchers modeling SRM activities is how long after deployment of geoengineering forcings it may be before one can determine the global

and regional effects with some confidence. As a part of the perturbed physics ensemble, I deployed a scenario ("B-test" in Figure 6.1) that makes small, periodic perturbations to stratospheric optical properties. This is consistent with work that Doug MacMynowski of Caltech and colleagues have done in which they simulated such periodic perturbations using HadCM3L with a single set of physical parameters (MacMynowski et al.). By setting up this scenario in a similar fashion to the ones in their experiment (e.g., using a sinusoidal forcing), the results should be able to complement their findings. As such, scenario B-test in Figure 2 shows a sinusoidal scenario with a period of two years and amplitude of about 0.04 units of optical depth (about 1 W/m2 of radiative forcing). The data from these simulations are yet to be analyzed

Scenarios for Rapid Cooling

While examining the climate response to scenarios designed to stabilize mean global temperature has been a helpful framework for examining relative sensitivities of regional climates to SRM and behavior of regional climates over time, long-term global temperature stabilization would not necessarily be the aim of a realistic SRM scheme. For instance, one situation in which SRM might be used is to force a rapid cool down of polar regions to stop the progress of ice melt and halt associated ice albedo feedbacks. The full set of PPE SRM scenarios included two in which SRM forcings are initiated later in the simulations (in model years 2020 and 2035, respectively) and slowly ramp up the level of SRM to and over the level in the temperature stabilization scenarios and then down again. ("C-ramp1" and "C-ramp2" in Figure 6.1)

Several SRM modeling studies have shown that if one were to use SRM to cool the climate and then abruptly discontinue it without a compensatory reduction in atmospheric greenhouse gases, rapid planetary warming would occur. Matthews and Caldeira (2007) Another scenario discontinues "injection" of stratospheric aerosols at their peak forcing

(resulting in a decay to essentially zero forcing within two years). ("D-cut" in Figure 6.1) This scenario is designed to provide information about potential "worst-case" disruptions to SRM as they might be implemented in the real world. Again, analysis of the results has not yet been completed

Additional Climate Modeling Experiments

Figure 3.9 showed one of the interesting results from the first CPDN SRM simulations, however due to the coarse resolution of the output of that experiment, and of the model itself, it is difficult to understand the underlying mechanisms driving this regional divergence or the significance of the result in terms of impacts on ecosystems, agriculture and other small scale systems. Such systems are affected by not just mean seasonal temperatures or precipitation, but precipitation variability and extremes, as well as the distribution of these effects among specific geographic areas. CPDN has recently started a regional climate modeling experiment using the global model, HadAM3P, coupled with HadRM3. HadAM3P hass higher resolution than HadCM3, but with prescribed sea surface temperatures. By embedding a regional climate model within a global model, one can observe the details of a climate change as it occurs in given regional system while still reaping the advantages of the dynamic interactions with the global system under anthropogenic climate change and under a given SRM forcing.

The regional modeling capacity of CPDN could be used for a complementary modeling experiment to be conducted as follow up to the PPE described above. The tool can, in theory, be applied to any region of interest. Currently, CPDN regional experiments are set up for Western North America, Southern Africa and Europe. It would be a natural first step in regional modeling of SRM activities to look at detailed regional responses in these areas, but the next step would be to set up regional modeling experiments for regions

where interesting regional effects were observed in the earlier experiments, such as those in South Asia and East Asia.

A point of particular contention in SRM research is whether these activities could result in the shutdown of the South Asian monsoon (Robock et al., 2008). Our results from the CPDN simulations of SRM to date do not show significant negative seasonal precipitation effects from SRM forcings in this region, but one goal of follow up simulations could be to clarify effects on South Asian precipitation and the anomalously different response of this region to SRM forcings.

Real Impacts Assessment

An SRM-modified world would exhibit a number of environmental differences that would have impacts on plants. Temperatures will be relatively cooler compared to the greenhouse gas-warmed world without SRM, precipitation will change (differently for different regions) and more diffuse sunlight and higher concentrations of carbon dioxide will both encourage plant growth. To date, there has been no published work that directly examines the effects this type of geoengineering would have on agricultural productivity or terrestrial ecosystems, impacts relevant to the decision making process of any nation or actor considering SRM.

Discussions of the effects of SRM have, to date, focused primarily on macroscopic changes to temperature and precipitation. A substantive comparison of the impacts of SRM to those of other climate change will require information about expected changes to key environmental systems in the context of socioeconomic indicators. Several recent papers have attempted to contextualize SRM modeling experiments in impacts-relevant ways (e.g., Moreno-Cruz et al. (2011) and Irvine et al. (2010)), but no one has yet taken the next step of using high-resolution climate model output to directly model the health and

productivity of natural and managed ecosystems, hydrologic systems and other important environmental indicators. The output from experiments conducted for this thesis could be adapted for such modeling exercises. Or, ideally, regional modeling experiments proposed above will be produced at spatial and temporal resolutions appropriate for feeding into regional models of such systems. A natural extension of the project would be to use these data to examine some of the true tradeoffs one may encounter when making decisions about future use of SRM.

In a changing climate of any form, a priority must remain food and water security along with avoiding dangerous impacts, many of which are associated with the hydrological cycle (e.g. flood and drought risk). It is essential therefore to understand the influence of deliberate climate modification on such impacts. Such forward predictions can determine whether additional adaptation measures are needed in a world with rising greenhouse gases concurrent with geo-engineering. It might emerge that the "side-effects" of proposed climate modification on the hydrological cycle are such that only certain capped levels of solar radiation management should be contemplated. Diagnostics of surface rainfall from the CPDN simulations could be used to force a continuous rainfall-runoff description of land surface response.

Appendix A

climate*prediction*.net Relevant Model Parameters & Forcing Files

Table A.1: climate*prediction*.net"Perturbable"Parameters and Forcing Files. See results.cpdn.org for more information.

Parameter or Forcing File	Descripiton	Default Value
alpham	albedo at melting point of ice	0.5
asym lambda	neutral mixing length	0.15
charnock	roughness lengths over the sea	0.012
cloudtau	time a circulating air parcel remains in a cloud	1.08E+04
ct	accretion constant	1.00E-04
cw land	preciptiation threshold over land	2.00E-04
cw sea	preciptiation threshold over sea	5.00E-05
diff coeff	horizontal diffusion coefficient	5.470e+08, 5.470e+08, 5.470e+08, 5.470e+08, 5.470e+08...
diff coeff q	horizontal diffusion coefficient for water vapour	5.470e+08, 5.470e+08, 5.470e+08, 5.470e+08, 5.470e+08...
diff exp	exponent of the horizontal diffusion	3, 3, 3, 3, 3...
diff exp q	exponent of the horizontal diffusion for water vapour	3, 3, 3, 3, 3...

Table A.1: climate*prediction*.net"Perturbable"Parameters and Forcing Files. See results.cpdn.org for more information.

Parameter or Forcing File	Descripiton	Default Value
dlat	IC ensemble: latitude position of dtheta perturbation	37
dlon	IC ensemble: longditude position of dtheta perturbation	48
dtheta	IC ensemble: initial condition parameter	0
dtice	temperature range of ice albedo variation	10
eacf	empirically adjusted cloud fraction	0.5, 0.5, 0.5, 0.5, 0.5 ...
eddydiff	slab ocean to ice heat transfer	3.75E-04
entcoef	entrainment coefficient	3
file atmos	atmospheric start file	yafbg.astart
file flux	Ocean: heat and salinity flux adjustment file particular to ocean spinup	1081 flux corr.anc
file ghg	greenhouse gas emissions scenario definition file	ghg a1b
file nh3so2	sulphate emissions file	NULL
file nh3so2 mod	file compression for sulphate emissions file	NULL
file nick	Ocean: heat and salinity flux adjustment file particular to atmospheric physics configuration	nick zeroflux.anc
file ocean	ocean start file	yafbg.ostart
file ozone	[undefined]	ozone hadcm3 1900
file ozone mod	[undefined]	NULL
file solar	solar v01	solar v01
file spec lw	spec3a lw 3 asol2c hadcm3	spec3a lw 3 asol2c hadcm3
file spec sw	spec3a sw 3 asol2b hadcm3	spec3a sw 3 asol2b hadcm3
file sulphate	[undefined]	NULL
file sulphox	[undefined]	NULL
file sulphox mod	[undefined]	NULL
file volc	natural volcanic emissions definition file	NAT VOLC

Table A.1: climate*prediction*.net"Perturbable"Parameters and Forcing Files. See results.cpdn.org for more information.

Parameter or Forcing File	Descripiton	Default Value
file volc mod	compression for natural volcanic emissions file	NULL
file volcanic	volcanic forcing scenario	volc v01
filtering safety factor	Filtering safety factor	NULL
go	stability dependence of turbulent mixing coefficients	10
haney	Ocean: Haney heat forcing coefficient	81.88
haneysfact	Ocean: Haney salinity forcing factor	0.25
i cnv ice lw	type for convective ice	1
i cnv ice sw	type for convective water	3
i st ice lw	type for stratiform ice	1
i st ice sw	type for stratiform water	2
ice size	ice size in radiation	0.00003
isopyc	Ocean: isopycnal diffusion of tracer at surface	1.00E+03
kay gwave	Surface stress constant for gravity wave drag	2.00E+04
kay lee gwave	Surface stress constant for lee gravity wave drag	3.00E+05
kay lee gwdrag	[undefined]	3.00E+05
lo	sulphate mass scavenging parameter L0	6.50E-05
l1	sulphate mass scavenging parameter L1	2.96E-05
lhaney	Ocean: Haney logical flag	FALSE
mldel	Ocean: decay of wind mixing energy with depth	1.00E+02
mllam	Ocean: wind mixing energy scaling factor	0.7
num star	threshold for condensation onto accumulation mode particles	1.00E+06
ocbohaney	[undefined]	1
ohsca	Sulphur cycle: Scaling parameter for OH field	

Table A.1: climate*prediction*.net"Perturbable"Parameters and Forcing Files. See results.cpdn.org for more information.

Parameter or Forcing File	Descripiton	Default Value
paramname	Parameter name (matches simulation name)	NULL
r layers	root depth	4, 4, 3, 3...
rhcrit	critical relative humidity	0.95, 0.90, 0.85, 0.70, 0.70...
sc	solar constant	NULL
so2 high level	Sulphur Cycle: Model level for SO2 (High level) emissions	3
start level gwdrag	lowest model level for application of gravity wave drag	3
umid		NULL
vdiffdepth	Ocean: increase of background vertical mixing of tracer with depth	2.80E-08
vdiffsurf	Ocean: background vertical mixing of tracer (diffusion) at surface	1.00E-05
vertvisc	Ocean: background vertical mixing of momentum (viscosity)	1.00E-05
vf1	ice fall speed	1
visbeck	Ocean: Visbeck scheme	FALSE
volsca	Sulphur cycle: Scaling factor for emission from natural (volcanic) emissions	1
zofsea	surface fluxes over tropical oceans	1.30E-03

Appendix B

Notes on Regional Normalizations

The primary normalization scheme used to comparing the anomalies observed in the 23 macro-regions analyzed is seasonal interannual variability from the baseline dataset (1990s transient standard physics simulations). This normalization scheme was selected as a recent common reference point that could provide a metric for how the general changes that happened over time in the SRM and no-SRM simulations compared to the (simulated) mean climate state and year-to-year variability of each region in the recent past with the logic that from an impacts relevance standpoint, the smaller changes to the mean climate of a region are compared to interannual variability in the recent past, the easier such changes will be to adapt to. Other approaches to normalization can yield somewhat different results. One standard method of displaying the significance of decadal climate anomalies is to normalize them by the decadal variability of a long control run of the model used. For comparison, we show in Figure , a version of Figure , but for annual means. On top (a), results are normalized as in Figure ; and on the bottom (b), those same annual means are normalized using decadal internal variability from a 116-decade control run of HadCM3 (with no perturbations to greenhouse gases, solar variability, etc.)

Unsurprisingly, considering the large perturbations we make to longwave and shortwave forcings in our transient scenarios, compared with decadal variability of the control climate, the decadal average changes observed between different levels of SRM appear even more significant. The green horizontal lines in panel (b) show the 5 standard deviations boundary. Rarely is there an SRM scenario that returns a regional climate to within one standard deviation of its baseline for both temperature and precipitation As Figure shows, there are significant differences in the relative sensitivities of various regions to SRM forcings. However, using this alternative approach to normalizing regional results does not affect our main findings that: (1) the "optimal" level of SRM varies by region and (2) the regional climates diverge disparately from baseline climate states over time with continued global-mean surface temperature stabilization under rising atmospheric CO2.

Further this alternative normalization does not impact the results presented in Figures , and in any way, nor the character of the result in Figure , nor any of the qualitative discussion of regional results presented in the main text. Estimates of surface tempera-

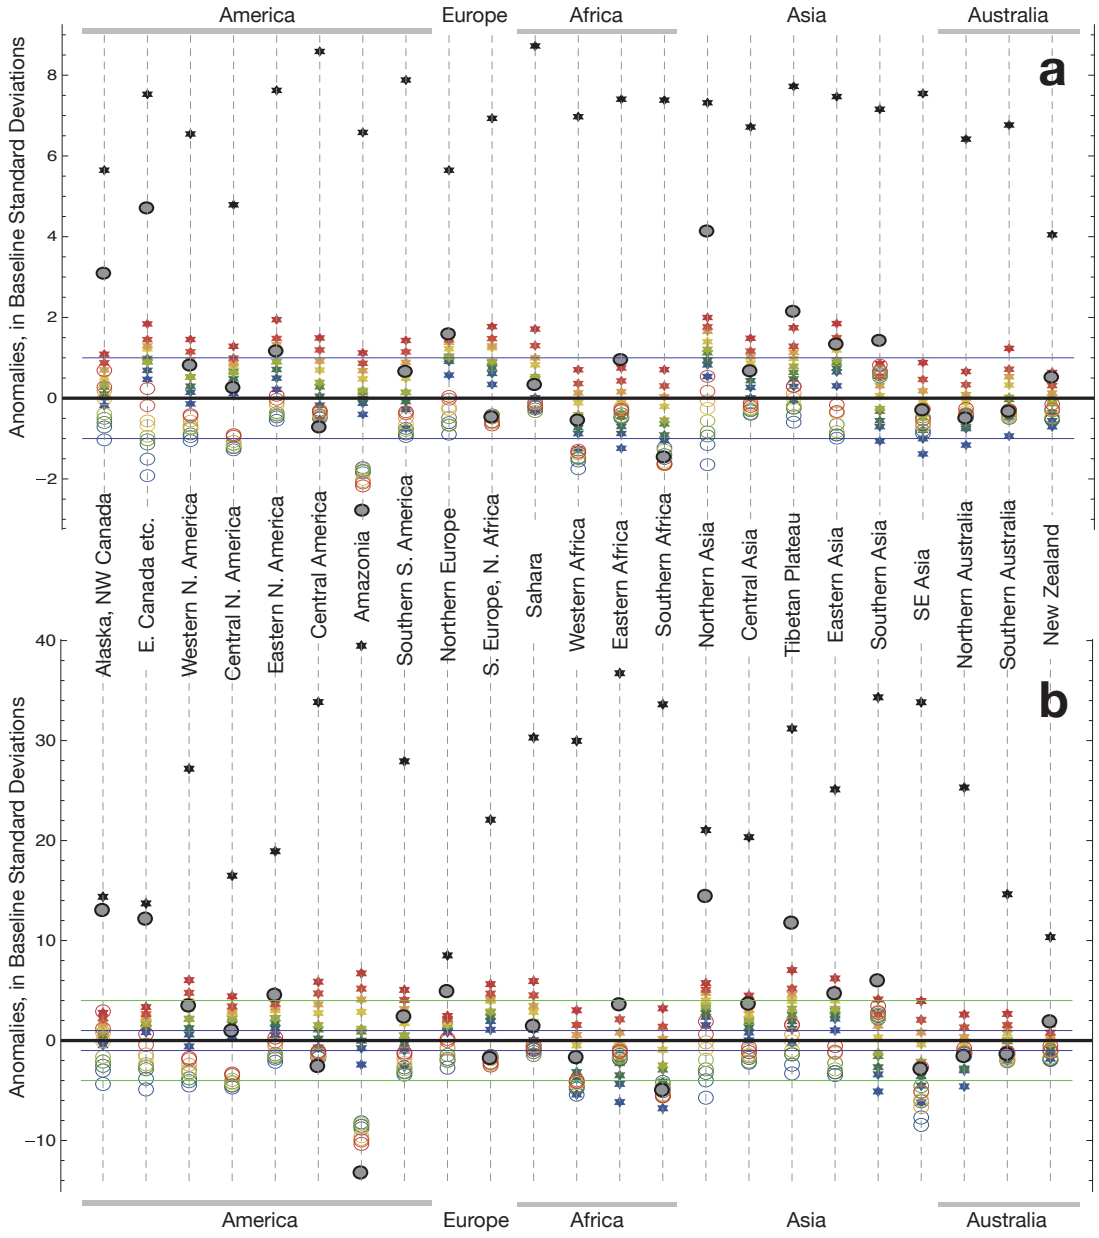

Figure S6. Normalized Regional Temperature and Precipitation Anomalies computed as the difference between average values for the decade of the 2070s and the 1990s in units of **(a)** baseline interannual standard deviations and **(b)** control-run interdecadal standard deviations. Stars show mean temperature data for SRM (colours) and no-SRM (black) scenarios. Open circles show mean precipitation data. Blue horizontal lines on each plot show the cut off for one standard deviation. Green horizontal lines in panel (b) show the cut off for five standard deviations. Each coloured point represents data from 60 simulations averaging over 6 SRM scenarios. Each black point represents data from 30 no-SRM simulations. The 1990s baseline dataset is compiled from 30 standard physics simulations. The control climate dataset comes from a 116-decade segment of a long control run of HadCM3.

ture variability HadCM3 have been shown to be in close agreement with observed climate data (Collins et al., 2001), though it has also been shown to under-represent precipitation variability (Lambert et al., 2004). All AOGCM's have limitations in their ability to accurately model regional precipitation, even absent any debate over normalization. For this reason we have avoided drawing quantitative conclusions about regional responses to SRM forcings.

Appendix C

Experiment 1: Test of Linearity of Regional Responses

A strong linear response of regional temperature and precipitation to changes in stratospheric aerosol optical depth (AOD – i.e., the amount of SRM) in the forcing range of interest is important because it means that in future experiments (i.e., Experiment 2), far fewer SRM scenarios need be simulated in order to get a clear picture of regional sensitivities.This appendix presents tables with the regression coefficients and $\bar{R}^2$ values associated with annual, Northern Hemisphere summer (JJA) and Northern Hemisphere winter (DJF) decadal mean temperature and precipitation anomalies for the 2020s and the 2070s.

Table C.1: Test of linearity on Radiative Forcing for Decadal-Mean, Annual-Mean Regional Temperatures.

	2020s							2070s						
Region	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$
ala	1.78	-0.0019	0.96	2.12	-0.0027	4.38E-07	0.96	4.62	-0.0020	0.87	6.90	-0.0042	5.06E-07	0.87
cgi	1.85	-0.0017	0.98	1.88	-0.0018	4.87E-08	0.98	3.56	-0.0014	0.84	7.67	-0.0053	9.09E-07	0.85
wna	1.29	-0.0014	0.98	1.36	-0.0016	7.93E-08	0.98	3.69	-0.0016	0.92	4.92	-0.0027	2.73E-07	0.92
cna	1.6	-0.0017	0.97	1.71	-0.0019	1.39E-07	0.97	4.53	-0.0018	0.91	4.07	-0.0014	-1.02E-07	0.90
ena	1.14	-0.0012	0.98	1.19	-0.0013	6.90E-08	0.98	3.04	-0.0012	0.92	4.14	-0.0023	2.43E-07	0.92
cam	1.02	-0.0011	0.99	1.08	-0.0013	7.17E-08	0.99	3.04	-0.0014	0.96	3.38	-0.0017	7.34E-08	0.95
amz	1.25	-0.0014	0.98	1.41	-0.0018	2.11E-07	0.98	3.90	-0.0017	0.93	5.64	-0.0034	3.86E-07	0.93
ssa	0.83	-0.0009	0.98	0.91	-0.0011	1.08E-07	0.98	2.60	-0.0012	0.93	2.88	-0.0014	6.12E-08	0.93
neu	1.52	-0.0014	0.95	1.42	-0.0011	-1.33E-07	0.95	3.31	-0.0012	0.74	3.91	-0.0018	1.32E-07	0.73
seu	1.34	-0.0013	0.98	1.36	-0.0014	2.01E-08	0.98	3.35	-0.0013	0.92	4.29	-0.0022	2.09E-07	0.92
sah	1.38	-0.0014	0.99	1.46	-0.0016	1.08E-07	0.99	3.51	-0.0015	0.96	5.03	-0.0030	3.36E-07	0.96
waf	0.81	-0.0011	0.98	0.86	-0.0012	6.66E-08	0.98	2.79	-0.0014	0.94	3.38	-0.0019	1.31E-07	0.93
eaf	0.97	-0.0012	0.99	1.07	-0.0015	1.31E-07	0.99	3.12	-0.0015	0.94	4.02	-0.0024	1.99E-07	0.94
saf	0.92	-0.0011	0.99	1	-0.0013	9.70E-08	0.99	2.91	-0.0014	0.95	3.83	-0.0023	2.03E-07	0.94
nas	1.73	-0.0017	0.98	1.62	-0.0014	-1.38E-07	0.98	4.54	-0.0018	0.94	4.69	-0.0019	3.18E-08	0.94
cas	1.54	-0.0016	0.98	1.53	-0.0016	-2.16E-08	0.98	3.84	-0.0016	0.94	3.76	-0.0015	-1.65E-08	0.94
tib	1.56	-0.0017	0.99	1.76	-0.0022	2.58E-07	0.99	4.14	-0.0017	0.94	5.08	-0.0026	2.06E-07	0.94
eas	1.16	-0.0012	0.99	1.15	-0.0012	-7.60E-09	0.99	3.18	-0.0013	0.92	4.72	-0.0027	3.39E-07	0.93
sas	0.87	-0.0011	0.98	0.97	-0.0014	1.36E-07	0.98	2.83	-0.0013	0.94	3.66	-0.0021	1.84E-07	0.94
sea	0.75	-0.001	0.99	0.83	-0.0011	1.01E-07	0.99	2.39	-0.0012	0.95	3.07	-0.0018	1.50E-07	0.95
nau	0.89	-0.0011	0.98	0.91	-0.0011	2.01E-08	0.98	2.80	-0.0014	0.95	3.74	-0.0023	2.09E-07	0.95
sau	0.71	-0.0007	0.96	0.69	-0.0006	-3.60E-08	0.95	2.29	-0.0011	0.93	2.71	-0.0015	9.34E-08	0.93
nz	0.74	-0.0008	0.94	0.71	-0.0007	-3.90E-08	0.94	2.11	-0.0010	0.87	2.81	-0.0017	1.56E-07	0.87

Table C.2: Test of linearity on Radiative Forcing for Decadal-Mean, Seasonal-Mean (JJA) Regional Temperatures.

	2020s							2070s						
Region	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$
ala	1.36	-0.0013	0.96	1.59	-0.0019	3.07E-07	0.96	3.46	-0.0014	0.85	6.62	-0.0044	7.00E-07	0.85
cgi	1.39	-0.0012	0.98	1.41	-0.0012	2.89E-08	0.98	2.51	-0.0009	0.77	6.33	-0.0045	8.46E-07	0.79
wna	1.33	-0.0014	0.99	1.49	-0.0018	2.08E-07	0.99	3.69	-0.0015	0.94	6.63	-0.0043	6.50E-07	0.95
cna	1.61	-0.0017	0.95	1.92	-0.0025	4.02E-07	0.96	4.51	-0.0018	0.86	4.20	-0.0015	-6.88E-08	0.86
ena	1.13	-0.0012	0.98	1.17	-0.0012	4.69E-08	0.98	2.94	-0.0012	0.92	2.95	-0.0012	3.69E-09	0.91
cam	0.98	-0.0012	0.98	1.11	-0.0015	1.71E-07	0.98	3.06	-0.0014	0.94	3.32	-0.0016	5.60E-08	0.94
amz	1.20	-0.0013	0.97	1.34	-0.0016	1.85E-07	0.97	3.80	-0.0017	0.92	5.44	-0.0032	3.64E-07	0.92
ssa	0.77	-0.0008	0.97	0.88	-0.0011	1.47E-07	0.97	2.45	-0.0011	0.90	2.59	-0.0012	3.07E-08	0.90
neu	1.38	-0.0012	0.95	1.19	-0.0008	-2.41E-07	0.95	2.91	-0.0010	0.73	4.08	-0.0021	2.59E-07	0.72
seu	1.42	-0.0014	0.98	1.40	-0.0014	-2.70E-08	0.98	3.61	-0.0014	0.92	3.83	-0.0016	4.74E-08	0.92
sah	1.28	-0.0014	0.99	1.35	-0.0016	9.13E-08	0.99	3.66	-0.0016	0.95	5.17	-0.0030	3.35E-07	0.95
waf	0.73	-0.0010	0.97	0.77	-0.0011	5.53E-08	0.97	2.72	-0.0014	0.93	2.92	-0.0016	4.52E-08	0.93
eaf	0.95	-0.0013	0.99	1.06	-0.0016	1.43E-07	0.99	3.24	-0.0016	0.94	4.37	-0.0027	2.50E-07	0.94
saf	1.00	-0.0013	0.98	1.04	-0.0013	4.61E-08	0.98	3.11	-0.0016	0.94	3.74	-0.0022	1.40E-07	0.94
nas	1.47	-0.0014	0.99	1.53	-0.0015	7.36E-08	0.99	3.89	-0.0015	0.95	4.23	-0.0018	7.70E-08	0.95
cas	1.53	-0.0016	0.97	1.52	-0.0016	-9.14E-09	0.97	4.02	-0.0016	0.92	5.49	-0.0030	3.25E-07	0.92
tib	1.36	-0.0014	0.97	1.49	-0.0017	1.66E-07	0.97	3.80	-0.0015	0.93	6.04	-0.0036	4.97E-07	0.94
eas	1.07	-0.0012	0.98	1.07	-0.0012	-4.12E-10	0.98	3.08	-0.0012	0.94	4.97	-0.0030	4.19E-07	0.94
sas	0.76	-0.0011	0.97	0.81	-0.0012	6.32E-08	0.97	2.67	-0.0014	0.94	2.67	-0.0014	1.34E-09	0.93
sea	0.74	-0.0010	0.99	0.80	-0.0011	7.58E-08	0.99	2.41	-0.0012	0.94	2.85	-0.0016	9.86E-08	0.94
nau	0.87	-0.0011	0.96	0.92	-0.0012	6.51E-08	0.95	2.48	-0.0012	0.88	3.97	-0.0027	3.31E-07	0.88
sau	0.71	-0.0007	0.95	0.70	-0.0006	-1.34E-08	0.95	2.10	-0.0010	0.90	3.31	-0.0021	2.68E-07	0.90
nz	0.71	-0.0008	0.92	0.78	-0.0010	1.01E-07	0.92	2.20	-0.0011	0.88	2.87	-0.0017	1.47E-07	0.88

Table C.3: Test of linearity on Radiative Forcing for Decadal-Mean, Seasonal-Mean (DJF) Regional Temperatures.

	2020s							2070s						
Region	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$
ala	2.37	-0.0026	0.94	2.90	-0.0038	6.93E-07	0.94	6.69	-0.0029	0.86	6.60	-0.0028	-1.86E-08	0.85
cgi	2.24	-0.0021	0.97	2.21	-0.0020	-4.00E-08	0.97	4.51	-0.0017	0.85	9.08	-0.0061	1.01E-06	0.86
wna	1.35	-0.0014	0.95	1.36	-0.0015	1.00E-08	0.95	3.82	-0.0016	0.83	3.15	-0.0010	-1.50E-07	0.82
cna	1.64	-0.0017	0.91	1.37	-0.0010	-3.50E-07	0.91	4.78	-0.0020	0.84	5.91	-0.0030	2.51E-07	0.84
ena	1.12	-0.0012	0.94	1.04	-0.0011	-9.80E-08	0.94	3.16	-0.0013	0.86	5.72	-0.0037	5.66E-07	0.87
cam	1.08	-0.0011	0.97	1.04	-0.0010	-5.25E-08	0.97	3.05	-0.0013	0.93	4.50	-0.0027	3.20E-07	0.93
amz	1.23	-0.0013	0.98	1.39	-0.0017	2.15E-07	0.98	3.80	-0.0017	0.93	6.11	-0.0039	5.11E-07	0.93
ssa	0.88	-0.0009	0.98	0.90	-0.0009	2.75E-08	0.98	2.65	-0.0012	0.94	3.47	-0.0019	1.82E-07	0.94
neu	1.57	-0.0015	0.90	1.65	-0.0017	1.02E-07	0.89	3.76	-0.0014	0.69	1.53	0.0007	-4.94E-07	0.68
seu	1.28	-0.0012	0.93	1.52	-0.0018	3.14E-07	0.93	3.12	-0.0012	0.89	5.40	-0.0034	5.05E-07	0.89
sah	1.50	-0.0014	0.97	1.60	-0.0017	1.29E-07	0.97	3.42	-0.0015	0.92	6.52	-0.0044	6.85E-07	0.92
waf	0.94	-0.0011	0.98	0.97	-0.0012	3.40E-08	0.98	2.90	-0.0013	0.92	3.77	-0.0022	1.92E-07	0.92
eaf	1.05	-0.0012	0.98	1.12	-0.0013	8.96E-08	0.98	3.10	-0.0014	0.92	3.95	-0.0022	1.88E-07	0.92
saf	0.93	-0.0011	0.99	1.01	-0.0013	1.12E-07	0.99	2.94	-0.0014	0.93	4.19	-0.0026	2.77E-07	0.93
nas	1.88	-0.0018	0.95	1.72	-0.0015	-2.04E-07	0.95	4.96	-0.0019	0.88	5.07	-0.0020	2.24E-08	0.88
cas	1.47	-0.0015	0.92	1.35	-0.0012	-1.58E-07	0.92	3.56	-0.0015	0.85	2.30	-0.0003	-2.80E-07	0.85
tib	1.71	-0.0019	0.96	1.86	-0.0022	1.95E-07	0.96	4.31	-0.0019	0.85	4.28	-0.0018	-6.98E-09	0.85
eas	1.28	-0.0013	0.97	1.27	-0.0013	-1.07E-08	0.97	3.27	-0.0013	0.85	4.40	-0.0024	2.51E-07	0.85
sas	0.94	-0.0012	0.95	1.14	-0.0017	2.55E-07	0.95	3.04	-0.0014	0.89	4.33	-0.0026	2.87E-07	0.89
sea	0.78	-0.0009	0.99	0.88	-0.0012	1.26E-07	0.99	2.38	-0.0011	0.94	3.12	-0.0018	1.64E-07	0.94
nau	0.89	-0.0010	0.96	0.92	-0.0011	3.50E-08	0.96	2.95	-0.0014	0.94	4.64	-0.0030	3.73E-07	0.94
sau	0.71	-0.0007	0.93	0.73	-0.0007	2.46E-08	0.93	2.48	-0.0011	0.92	3.49	-0.0021	2.25E-07	0.92
nz	0.71	-0.0007	0.89	0.62	-0.0005	-1.18E-07	0.89	1.94	-0.0009	0.77	1.62	-0.0006	-7.10E-08	0.77

Table C.4: Test of linearity on Radiative Forcing for Decadal-Mean, Annual-Mean Regional Precipitation Rates.

	2020s							2070s						
Region	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$
ala	0.12	-0.0001	0.91	0.12	-0.0001	-6.81E-10	0.91	0.34	-0.0002	0.84	0.76	-0.0006	9.47E-08	0.85
cgi	0.12	-0.0001	0.98	0.13	-0.0001	3.37E-09	0.98	0.22	-0.0001	0.87	0.26	-0.0001	7.25E-09	0.87
wna	0.03	-0.0001	0.63	0.01	0.0000	-2.32E-08	0.62	0.16	-0.0001	0.55	0.12	-0.0001	-8.66E-09	0.54
cna	0.00	-0.0001	0.28	-0.03	0.0000	-3.71E-08	0.28	0.06	-0.0001	0.31	0.14	-0.0002	1.68E-08	0.30
ena	0.10	-0.0001	0.53	0.10	-0.0001	9.47E-09	0.52	0.21	-0.0001	0.47	0.78	-0.0007	1.26E-07	0.48
cam	0.07	-0.0001	0.45	0.07	-0.0001	3.60E-10	0.44	0.03	0.0000	0.07	0.69	-0.0007	1.45E-07	0.09
amz	-0.24	0.0001	0.63	-0.27	0.0002	-3.91E-08	0.63	-0.68	0.0002	0.34	-1.56	0.0010	-1.93E-07	0.34
ssa	0.05	0.0000	0.40	0.06	0.0000	3.57E-09	0.39	0.10	-0.0001	0.35	0.37	-0.0003	5.79E-08	0.35
neu	0.07	-0.0001	0.76	0.05	0.0000	-2.70E-08	0.76	0.18	-0.0001	0.53	0.25	-0.0002	1.56E-08	0.52
seu	0.01	0.0000	0.08	0.02	0.0000	1.05E-08	0.06	-0.09	0.0000	0.06	0.04	-0.0001	3.00E-08	0.05
sah	0.02	0.0000	0.09	0.02	0.0000	4.56E-09	0.07	0.00	0.0000	-0.01	0.10	-0.0001	2.29E-08	-0.02
waf	-0.07	-0.0001	0.31	-0.15	0.0001	-1.08E-07	0.33	0.02	-0.0001	0.32	-0.32	0.0002	-7.59E-08	0.31
eaf	0.01	-0.0001	0.37	-0.03	0.0000	-4.18E-08	0.36	0.12	-0.0001	0.16	0.69	-0.0006	1.27E-07	0.16
saf	-0.14	0.0001	0.49	-0.14	0.0001	-1.33E-11	0.48	-0.24	0.0001	0.19	-0.09	-0.0001	3.19E-08	0.18
nas	0.09	-0.0001	0.97	0.08	-0.0001	-1.78E-08	0.97	0.24	-0.0001	0.88	0.25	-0.0001	2.99E-09	0.88
cas	0.05	0.0000	0.16	0.07	-0.0001	1.92E-08	0.15	0.07	0.0000	0.14	0.22	-0.0002	3.17E-08	0.13
tib	0.08	-0.0001	0.80	0.08	-0.0001	-8.87E-09	0.80	0.22	-0.0001	0.69	0.25	-0.0001	7.30E-09	0.69
eas	0.07	-0.0001	0.72	0.10	-0.0002	3.97E-08	0.72	0.33	-0.0002	0.56	0.25	-0.0001	-1.68E-08	0.56
sas	0.14	-0.0001	0.49	0.21	-0.0003	8.17E-08	0.49	0.29	-0.0001	0.16	1.30	-0.0010	2.24E-07	0.19
sea	0.17	-0.0003	0.51	0.05	0.0000	-1.53E-07	0.50	0.43	-0.0003	0.19	0.06	0.0000	-8.04E-08	0.18
nau	0.01	0.0000	0.01	0.17	-0.0004	2.03E-07	0.04	-0.01	0.0000	0.00	-0.58	0.0005	-1.27E-07	-0.02
sau	-0.02	0.0000	0.08	-0.03	0.0000	-4.05E-09	0.06	-0.01	0.0000	0.00	0.25	-0.0003	5.86E-08	-0.01
nz	0.09	-0.0001	0.51	0.10	-0.0001	1.24E-08	0.50	0.14	-0.0001	0.27	0.36	-0.0003	4.66E-08	0.26

Table C.5: Test of linearity on Radiative Forcing for Decadal-Mean, Seasonal-Mean (JJA) Regional Precipitation Rates.

	2020s							2070s						
Region	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$
ala	0.14	-0.0002	0.91	0.12	-0.0002	-1.93E-08	0.91	0.35	-0.0002	0.75	1.58	-0.0014	2.71E-07	0.80
cgi	0.08	-0.0001	0.91	0.11	-0.0002	3.67E-08	0.91	0.20	-0.0001	0.82	0.32	-0.0002	2.81E-08	0.82
wna	-0.01	0.0000	0.03	0.00	0.0000	2.94E-09	0.01	0.01	-0.0001	0.16	-0.14	0.0001	-3.40E-08	0.15
cna	-0.10	0.0001	0.07	-0.22	0.0004	-1.62E-07	0.09	-0.23	0.0000	-0.02	0.49	-0.0007	1.60E-07	-0.03
ena	0.09	0.0000	0.11	0.08	0.0000	-1.16E-08	0.09	0.09	-0.0001	0.11	0.05	0.0000	-8.07E-09	0.10
cam	0.13	-0.0001	0.17	0.08	0.0000	-6.84E-08	0.15	-0.06	0.0000	-0.02	2.27	-0.0022	5.16E-07	0.01
amz	-0.16	0.0001	0.40	-0.10	0.0000	7.33E-08	0.40	-0.62	0.0002	0.27	-0.88	0.0004	-5.78E-08	0.26
ssa	0.07	0.0000	0.21	0.07	0.0000	-9.71E-10	0.19	0.11	-0.0001	0.20	0.05	0.0000	-1.46E-08	0.19
neu	0.05	-0.0001	0.56	0.05	-0.0001	1.53E-09	0.55	0.00	0.0000	0.09	0.27	-0.0003	5.98E-08	0.08
seu	-0.04	0.0001	0.40	-0.05	0.0001	-1.31E-08	0.39	-0.13	0.0000	0.05	-0.08	0.0000	9.77E-09	0.03
sah	0.02	0.0000	0.18	0.03	0.0000	8.72E-09	0.17	0.03	0.0000	0.17	0.07	-0.0001	9.12E-09	0.15
waf	-0.14	-0.0001	0.32	-0.16	-0.0001	-3.60E-08	0.30	-0.02	-0.0002	0.32	-0.17	0.0000	-3.27E-08	0.30
eaf	-0.04	0.0000	0.00	-0.16	0.0003	-1.49E-07	0.04	0.03	0.0000	0.00	0.39	-0.0004	7.88E-08	-0.02
saf	-0.07	0.0000	0.13	-0.04	0.0000	3.32E-08	0.13	-0.13	0.0000	0.09	0.03	-0.0001	3.62E-08	0.08
nas	0.09	-0.0001	0.95	0.06	-0.0001	-4.25E-08	0.95	0.21	-0.0001	0.82	0.34	-0.0003	3.06E-08	0.82
cas	0.02	0.0000	-0.01	0.01	0.0000	-5.43E-09	-0.03	0.02	0.0000	0.03	0.06	-0.0001	7.49E-09	0.01
tib	0.12	-0.0001	0.76	0.13	-0.0001	1.26E-08	0.76	0.32	-0.0002	0.66	0.36	-0.0002	8.02E-09	0.65
eas	0.16	-0.0003	0.80	0.16	-0.0003	8.22E-10	0.80	0.60	-0.0003	0.67	-0.03	0.0003	-1.39E-07	0.67
sas	0.17	0.0000	-0.01	0.08	0.0002	-1.09E-07	-0.03	0.28	0.0000	-0.02	3.57	-0.0031	7.28E-07	0.02
sea	0.21	-0.0005	0.73	0.27	-0.0007	7.84E-08	0.73	0.63	-0.0005	0.36	-0.81	0.0008	-3.18E-07	0.35
nau	-0.01	0.0000	0.02	0.02	-0.0001	4.82E-08	0.01	-0.13	0.0000	-0.01	-1.92	0.0017	-3.98E-07	0.03
sau	-0.01	0.0000	-0.02	-0.07	0.0001	-7.90E-08	0.00	0.05	0.0000	0.05	0.51	-0.0005	1.02E-07	0.05
nz	0.11	-0.0002	0.45	0.08	-0.0001	-2.90E-08	0.44	0.31	-0.0002	0.22	-0.30	0.0004	-1.35E-07	0.21

Table C.6: Test of linearity on Radiative Forcing for Decadal-Mean, Seasonal-Mean (DJF) Regional Precipitation Rates.

	2020s							2070s						
Region	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$	α_1	β_1	$\overline{R^2}$	α_2	β_2	γ_2	$\overline{R^2}$
ala	0.11	-0.0001	0.67	0.15	-0.0002	4.88E-08	0.68	0.31	-0.0001	0.62	-0.05	0.0002	-7.92E-08	0.62
cgi	0.17	-0.0002	0.92	0.13	-0.0001	-5.45E-08	0.93	0.30	-0.0001	0.68	0.13	0.0000	-3.70E-08	0.68
wna	0.06	-0.0001	0.47	0.04	-0.0001	-2.12E-08	0.46	0.29	-0.0001	0.40	0.94	-0.0008	1.42E-07	0.40
cna	0.04	-0.0001	0.39	-0.07	0.0002	-1.46E-07	0.42	0.24	-0.0002	0.41	1.96	-0.0018	3.80E-07	0.47
ena	0.11	-0.0001	0.40	0.19	-0.0003	9.17E-08	0.40	0.40	-0.0002	0.36	1.31	-0.0011	2.03E-07	0.37
cam	0.09	-0.0001	0.05	0.05	0.0000	-5.12E-08	0.03	0.24	-0.0001	0.10	1.56	-0.0014	2.91E-07	0.11
amz	-0.27	0.0001	0.37	-0.32	0.0002	-6.74E-08	0.37	-0.65	0.0001	0.17	-2.75	0.0021	-4.66E-07	0.25
ssa	0.07	-0.0001	0.56	0.08	-0.0001	9.56E-09	0.56	0.21	-0.0001	0.45	1.00	-0.0009	1.75E-07	0.48
neu	0.10	-0.0002	0.67	0.15	-0.0003	5.96E-08	0.67	0.36	-0.0002	0.50	-0.23	0.0004	-1.31E-07	0.50
seu	0.05	0.0000	-0.02	0.00	0.0001	-6.39E-08	-0.02	-0.05	0.0000	0.00	0.86	-0.0008	2.01E-07	0.03
sah	-0.01	0.0000	0.09	-0.02	0.0000	-6.90E-09	0.07	-0.06	0.0000	0.13	0.07	-0.0001	2.73E-08	0.11
waf	-0.05	-0.0001	0.14	-0.10	0.0000	-6.19E-08	0.13	0.11	-0.0001	0.15	0.13	-0.0002	5.79E-09	0.14
eaf	-0.04	-0.0001	0.18	-0.10	0.0000	-8.33E-08	0.17	0.01	-0.0001	0.03	1.00	-0.0010	2.18E-07	0.02
saf	-0.22	0.0001	0.18	-0.26	0.0002	-5.58E-08	0.17	-0.32	0.0001	0.06	-0.48	0.0002	-3.58E-08	0.04
nas	0.11	-0.0001	0.92	0.10	-0.0001	-1.09E-08	0.92	0.30	-0.0001	0.85	0.30	-0.0001	-7.93E-10	0.85
cas	0.06	-0.0001	0.19	0.07	-0.0001	2.03E-08	0.17	0.09	-0.0001	0.10	0.00	0.0000	-2.06E-08	0.08
tib	0.03	-0.0001	0.24	0.03	0.0000	-7.66E-09	0.23	0.07	0.0000	0.09	-0.04	0.0001	-2.57E-08	0.08
eas	-0.02	0.0000	0.02	0.12	-0.0004	1.80E-07	0.11	-0.01	0.0000	0.00	0.18	-0.0002	4.17E-08	-0.01
sas	0.03	-0.0001	0.09	0.19	-0.0005	2.09E-07	0.11	0.02	0.0000	-0.01	-0.12	0.0001	-3.20E-08	-0.03
sea	0.27	-0.0003	0.23	0.05	0.0003	-2.89E-07	0.23	0.61	-0.0003	0.10	-0.20	0.0005	-1.78E-07	0.08
nau	0.12	-0.0001	0.01	0.32	-0.0005	2.55E-07	0.01	0.03	0.0000	-0.02	-0.98	0.0009	-2.25E-07	-0.03
sau	-0.04	0.0000	-0.02	-0.01	-0.0001	3.41E-08	-0.02	-0.02	0.0000	0.03	-0.41	0.0003	-8.64E-08	0.03
nz	0.04	0.0000	0.02	0.05	-0.0001	8.14E-09	0.00	0.06	-0.0001	0.04	-0.14	0.0001	-4.55E-08	0.03

Appendix D

Linearity Test from Residual Climate Response Model Paper

The design of the modeling experiment we obtained our data from, which tested the effects of many levels of SRM forcing [see Ricke et.al. 2010], allows us to test our linearity assumption explicitly. Figure D.1 presents the behavior of the data and tests for linearity of temperature and precipitation changes as a function of compensated radiative forcing. In all panels, the horizontal axis shows the level of SRM minus the mean natural volcanic aerosol forcing measured in Wm^{-2}. The dots represent 54 different levels of SRM and the colors show the effects for each of the 22 regions defined in the paper. The left-hand panel shows the results for temperature changes in 2030 measured as number of standard deviations away from the regional baseline with SRM implemented. This panel also shows high level of co-linearity across regions. The right-hand panel shows the results for compensated precipitation changes. We can see a clear correlation between precipitation and SRM forcing in each region, albeit noisier and less consistently linear than the temperature data. Regions are less collinear in terms of precipitation.

Tables D.1 and D.2 below show the regression results for temperature and precipitation, respectively. For both temperature and precipitation we first show a linear model fit. The first column shows the slope of the linear fit with the slope error in parenthesis, and the second column shows the fraction of variability explained by the linear model. The units of the slopes are $(Wm^{-2})^{-1}$. The linear fits explain more than 95% of the variability of the temperature data with coefficient error below 1%. In terms of precipitation, as it can also be seen in the figure above, the data is much noisier and the linear fit has a lower explanatory power. Regions for which the variance explained is low are those with a smaller slope coefficient. However, the coefficient error for each region is below 1%, and, although we do not report it on the table, all P-values are asymptotically equal to zero. The results of the quadratic fit show that the fraction of explained variability does not improve substantially compared to the linear model.

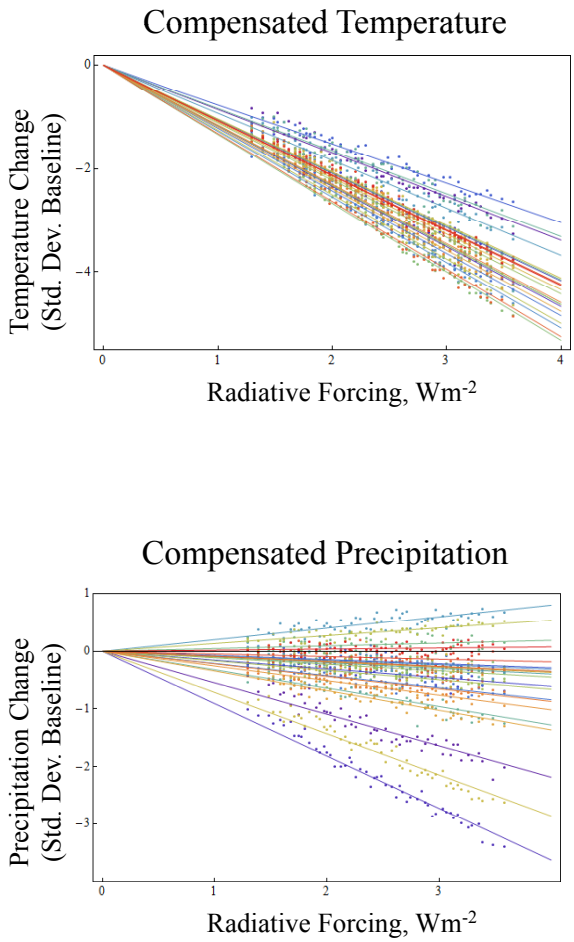

Figure D.1: Test of linearity of the compensated changes in precipitation and temperature.

Table D.1: Test of linearity on Radiative Forcing (RF): Temperature.

	Linear Fit=$a * RF$		Quadratic Fit=$a * RF + b * RF^2$		
	a (error)	Explained Variance	a (error)	b (error)	Explained Variance
Alaska, NW Canada	-0.85 (0.008)	95%	-0.70 (0.033)	-0.05 (0.012)	96%
E Canada etc.	-1.17 (0.007)	97%	-1.07 (0.035)	-0.04 (0.013)	97%
Western N America	-1.04 (0.005)	98%	-0.95 (0.024)	-0.04 (0.008)	98%
Central N America	-0.76 (0.006)	95%	-0.74 (0.029)	-0.01 (0.01)	95%
Eastern N America	-1.22 (0.007)	97%	-1.11 (0.034)	-0.04 (0.012)	98%
Central America	-1.27 (0.008)	97%	-1.07 (0.029)	-0.07 (0.011)	99%
Amazonia	-0.92 (0.006)	97%	-0.86 (0.03)	-0.02 (0.011)	97%
Southern S America	-1.16 (0.008)	97%	-1.00 (0.034)	-0.06 (0.012)	98%
Northern Europe	-0.83 (0.008)	93%	-0.65 (0.035)	-0.06 (0.013)	96%
S Europe, N Africa	-1.04 (0.006)	97%	-0.90 (0.024)	-0.05 (0.009)	98%
Sahara	-1.34 (0.006)	99%	-1.19 (0.021)	-0.05 (0.008)	99%
Western Africa	-1.11 (0.007)	97%	-0.98 (0.033)	-0.05 (0.012)	98%
Eastern Africa	-1.16 (0.008)	97%	-1.01 (0.033)	-0.06 (0.012)	98%
Southern Africa	-1.25 (0.007)	98%	-1.09 (0.027)	-0.06 (0.01)	99%
Northern Asia	-1.03 (0.008)	96%	-0.85 (0.03)	-0.07 (0.011)	98%
Central Asia	-1.09 (0.005)	98%	-1.00 (0.023)	-0.03 (0.008)	99%
Tibetan Plateau	-1.19 (0.006)	98%	-1.10 (0.026)	-0.03 (0.009)	99%
Eastern Asia	-1.06 (0.005)	98%	-0.97 (0.021)	-0.03 (0.008)	99%
Southern Asia	-1.15 (0.008)	97%	-0.98 (0.031)	-0.06 (0.011)	98%
Southeast Asia	-1.32 (0.008)	98%	-1.12 (0.026)	-0.07 (0.009)	99%
Northern Australia	-1.06 (0.006)	98%	-0.97 (0.026)	-0.03 (0.009)	98%
Southern Australia	-1.07 (0.007)	97%	-1.02 (0.032)	-0.02 (0.012)	97%

Table D.2: Test of linearity on Radiative Forcing (RF): Precipitation.

	Linear Fit=$a * RF$		Quadratic Fit=$a * RF + b * RF^2$		
	a (error)	Explained Variance	a (error)	b (error)	Explained Variance
Alaska, NW Canada	-0.55 (0.007)	89%	-0.42 (0.034)	-0.05 (0.012)	92%
E Canada etc.	-0.91 (0.007)	96%	-0.75 (0.027)	-0.06 (0.01)	98%
Western N America	-0.15 (0.006)	40%	-0.14 (0.031)	0.00 (0.011)	40%
Central N America	-0.07 (0.005)	28%	-0.02 (0.026)	-0.02 (0.009)	32%
Eastern N America	-0.21 (0.006)	60%	-0.22 (0.028)	0.00 (0.01)	60%
Central America	-0.08 (0.006)	41%	0.03 (0.023)	-0.04 (0.008)	59%
Amazonia	0.20 (0.006)	49%	0.25 (0.028)	-0.02 (0.01)	51%
Southern S America	-0.10 (0.005)	42%	-0.06 (0.024)	-0.02 (0.009)	45%
Northern Europe	-0.32 (0.007)	77%	-0.22 (0.032)	-0.03 (0.011)	81%
S Europe, N Africa	0.05 (0.006)	0%	0.07 (0.03)	-0.01 (0.011)	1%
Sahara	-0.09 (0.006)	19%	-0.10 (0.028)	0.00 (0.01)	20%
Western Africa	-0.11 (0.006)	39%	-0.06 (0.031)	-0.02 (0.011)	42%
Eastern Africa	-0.16 (0.006)	39%	-0.19 (0.029)	0.01 (0.01)	40%
Southern Africa	0.14 (0.004)	56%	0.11 (0.022)	0.01 (0.008)	57%
Northern Asia	-0.71 (0.007)	94%	-0.55 (0.026)	-0.06 (0.009)	97%
Central Asia	-0.09 (0.006)	21%	-0.10 (0.03)	0.00 (0.011)	21%
Tibetan Plateau	-0.34 (0.005)	85%	-0.29 (0.026)	-0.02 (0.009)	86%
Eastern Asia	-0.25 (0.005)	70%	-0.28 (0.025)	0.01 (0.009)	71%
Southern Asia	-0.22 (0.007)	19%	-0.32 (0.03)	0.04 (0.011)	34%
Southeast Asia	-0.09 (0.007)	34%	0.03 (0.032)	-0.04 (0.011)	49%
Northern Australia	-0.05 (0.008)	5%	-0.11 (0.037)	0.02 (0.013)	1%
Southern Australia	0.02 (0.004)	5%	0.00 (0.022)	0.01 (0.008)	6%

Appendix E

Supplemental Figures to Figure 4.4

Figure 4.4 shows an example of how regional responses to greenhouse gas and SRM forcings vary between models. The regions displayed were selected for simplicity sake. It is important to note that, unlike the Eastern North America and the Mediterranean, for some regions it is not possible to return temperature and precipitation to their baseline values even in the standard physics model. In addition, not all model versions show the clear distinct between the standard physics version and the highest response version. Figures E.1 and E.2 illustrate this point by showing the same data for each region, but with all model versions; and for each model version, but with all regions.

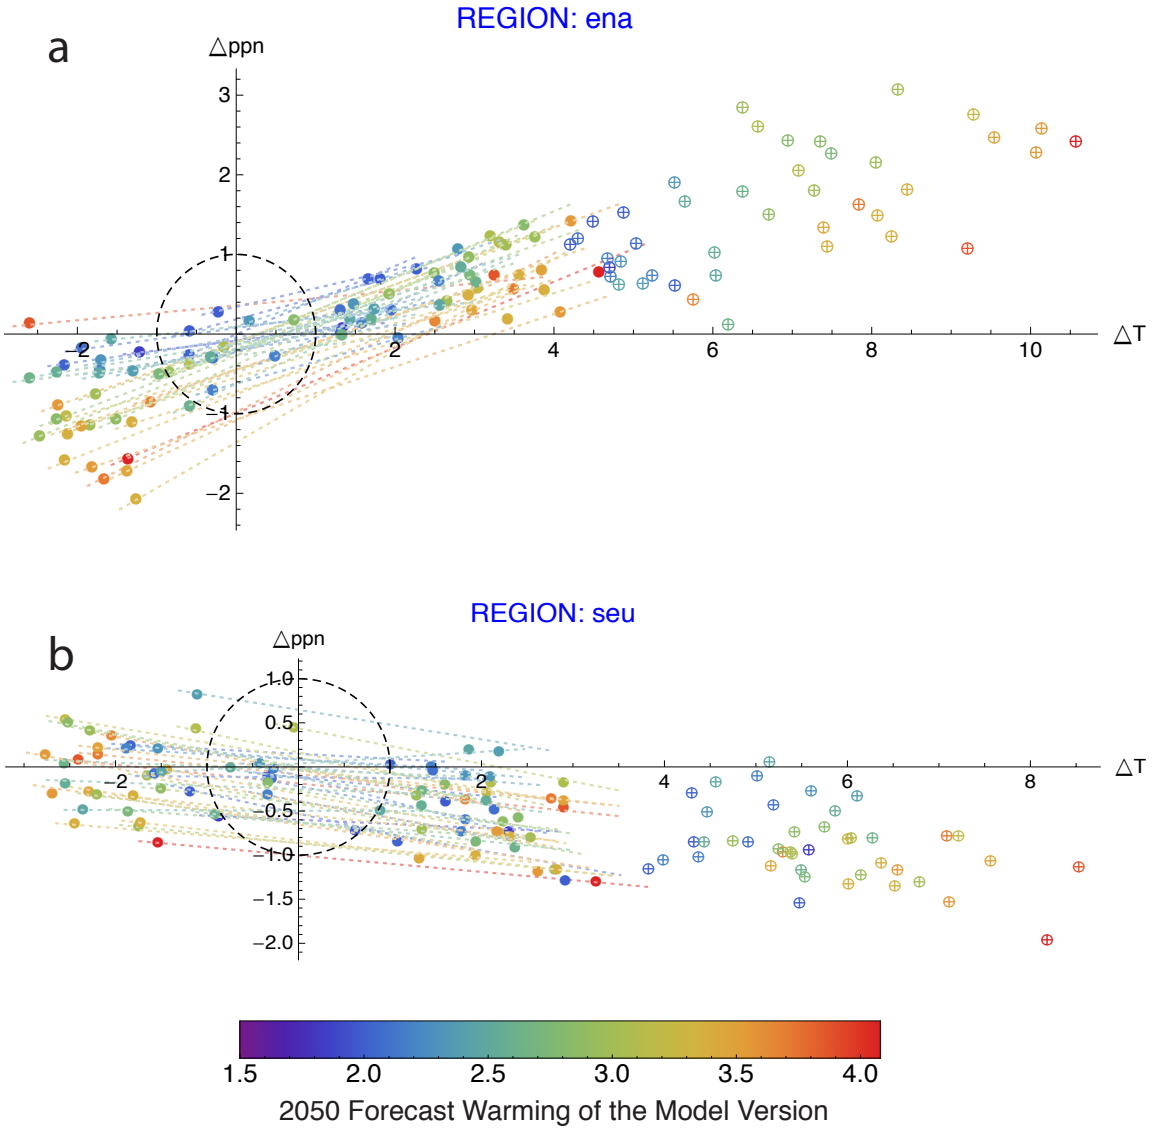

Figure E.1: Example of how regional responses to greenhouse gas and SRM forcings vary between models. The series of points show the no-SRM, low SRM and high SRM responses for the all model versions (shaded according to the key at the bottom for (a) Eastern North America, and (b) Southern Europe/ Northern Africa (referred to as Mediterranean in the main text). Changes in temperature and precipitation are calculated as the ten-year average 2050 minus 2000, divided by the interannual variability of the control climate (calculated with simulations that used identical paramenter combinations but no anthropogenic forcings of any kind). Dashed lines indicate the linear trajectory as SRM is increased or decreased.

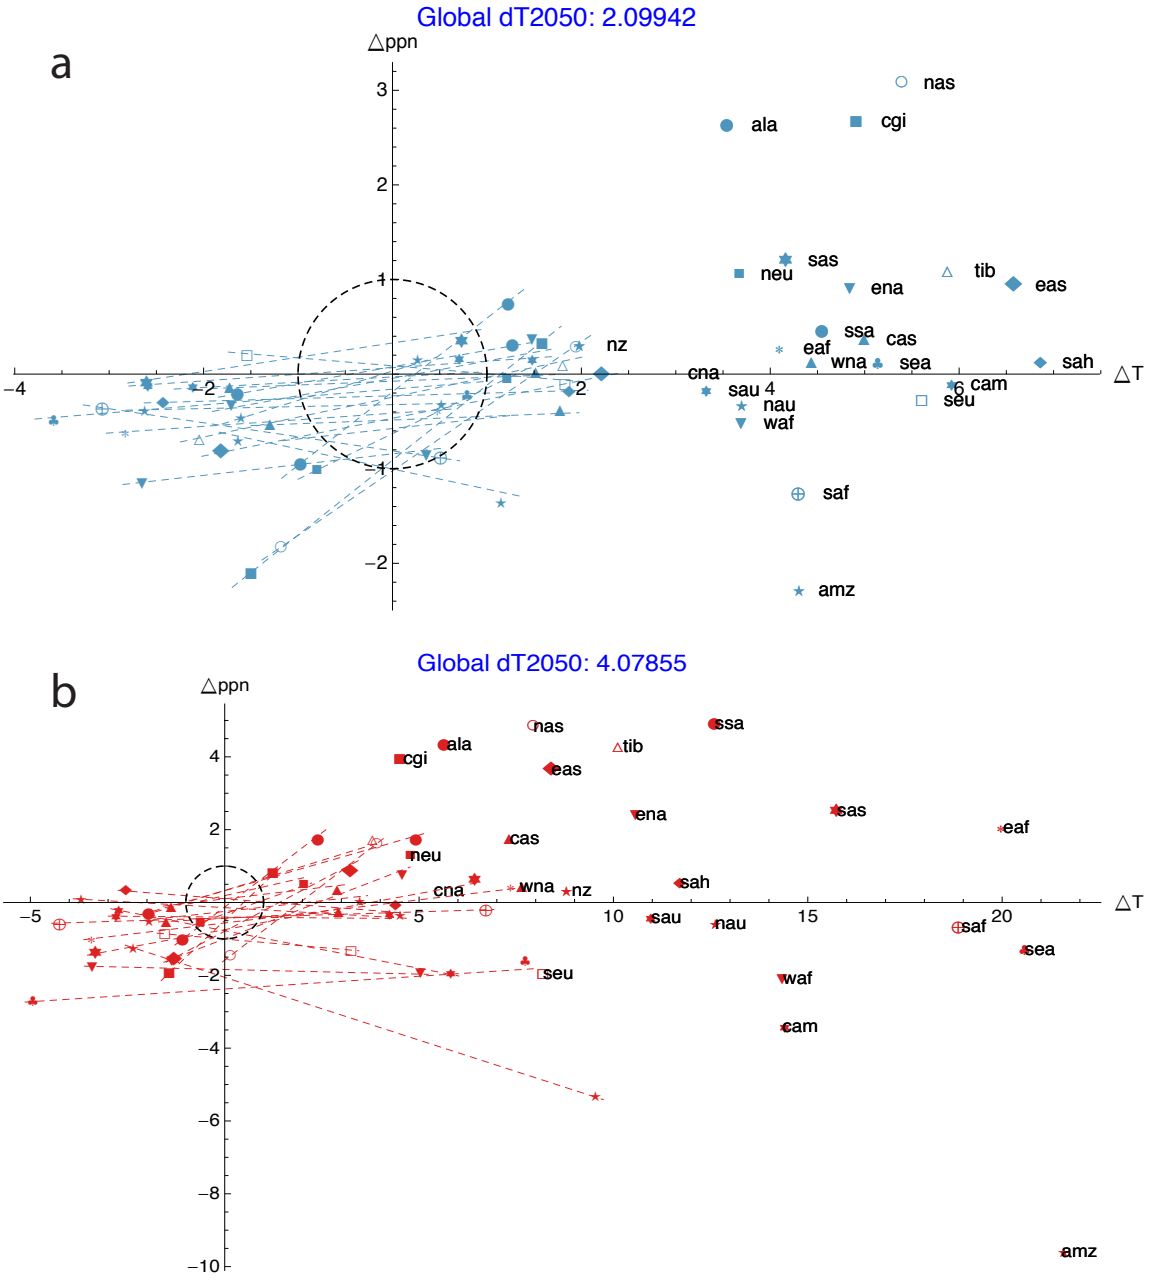

Figure E.2: Example of how regional responses to greenhouse gas and SRM forcings vary between models. The series of points show the no-SRM, low SRM and high SRM responses for all regions for (a) the standard settings of model physics parameters (ΔT_{2050}=2.1°C), and (b) our ensemble's highest temperature response model configuration ΔT_{2050}=4.1°C). Changes in temperature and precipitation are calculated as the ten-year average 2050 minus 2000, divided by the interannual variability of the control climate (calculated with simulations that used identical paramenter combinations but no anthropogenic forcings of any kind). Dashed lines indicate the linear trajectory as SRM is increased or decreased.

Appendix F

SRM Efficacy Metrics

The analysis presented in Chapter 4 relies on several measures designed to gauge how well SRM could work, in compensating for the effects of global warming in a single region and how well it works to restore regional climates overall and relative to no-SRM, once a particular level of SRM is set. This appendix provides the details of the derivation of these metrics.

These metrics are a tool to explicitly account for the precipitation and temperature tradeoffs that occur when attempting to stabilize regional climate changes using SRM. Precipitation and temperature changes only two, albeit very important, of the many variables likely to have climate related impacts. They are not intended as definitive normative measures of regional impacts or preferences. For example, when temperature rises in a given region, an increase in precipitation is required to maintain soil moisture. Any region would likely prefer a shift in their mean climate state of an increase of half a standard deviation in temperature with a half standard deviation more precipitation rather than an increase of half a standard deviation in temperature with half a standard deviation less precipitation. However, incremental changes to the amount of SRM never result in such a shift in precipitation space only. These metrics are specific to assessing SRM efficacy, but are obviously less informative with respect to a number of other alternative policies for dealing with climate change.

First, we consider OD*, the amount optical depth modification that returns a given regional climate closest to its baseline state in units of standard deviations, to measure the diversity of regional preferences for the amount of SRM. OD* is derived for a given region by first deriving the linear relationship between the specified value of stratospheric aerosol optical depth (od) and regional temperature (T) or precipitation (P) response. The data points used for these derivations were 10-year, initial condition ensemble means presented in units of standard deviations (regional interannual variability of the control (unperturbed) climate). So for a given region, time period (and season or annual-mean), there exist the following data points for the four SRM scenarios:

$$\begin{pmatrix} od_{noSRM} & T_{noSRM} & P_{noSRM} \\ od_{loSRM} & T_{loSRM} & P_{loSRM} \\ od_{medSRM} & T_{medSRM} & P_{medSRM} \\ od_{hiSRM} & T_{hiSRM} & P_{hiSRM} \end{pmatrix}$$

From the analysis of Experiment 1 output, we know that with sufficiently large initial condition ensembles, even with a small number of observations of regional responses to different levels of od, a linear fit works quite well. So using least squares regression, we derive:

$$T_{region}(od) = \alpha_{T_{reg}} + \beta_{T_{reg}} \cdot od$$
$$P_{region}(od) = \alpha_{P_{reg}} + \beta_{P_{reg}} \cdot od$$

To calculate OD* for a given region, then:

$$\textbf{Min}\ (\alpha_{T_{reg}} + \beta_{T_{reg}} \cdot od)^2 + (\alpha_{P_{reg}} + \beta_{P_{reg}} \cdot od)^2$$

so:

$$OD^*_{region} = \frac{-(\alpha_{T_{reg}} \cdot \beta_{T_{reg}} + \alpha_{P_{reg}} \cdot \beta_{P_{reg}})}{\beta^2_{T_{reg}} + \beta^2_{P_{reg}}}$$

Next, we consider regional anomalies, i.e., the regional temperature, precipitation or net temperature and precipitation changes, associated with a set amount of SRM to gauge likely regional satisfaction with and potential inequities in SRM's compensatory power relative to their baseline climate state. The amount of SRM is "democratically selected": simply the mean value of OD* for all regions, either unweighted (each region gets a vote), or weighted by regional population or economic output (each person or dollar gets a vote). (See Chapter 2 for information on the socioeconomic datasets used.) As such for a given weighting, w:

$$\overline{OD^*_w} = \frac{\sum_{r=1}^{23} w_{reg} \cdot OD^*_{reg}}{\sum_{r=1}^{23} w_{reg}}$$

and the regional climate anomalies (RA) associated with $\overline{OD^*_w}$ are:

$$\textbf{RA}_{reg} = (\alpha_{T_{reg}} + \beta_{T_{reg}} \cdot \overline{OD^*_w})^2 + (\alpha_{P_{reg}} + \beta_{P_{reg}} \cdot \overline{OD^*_w})^2$$

and the (weighted-)mean regional anomalies, as displayed in Figure 4.7 are :

$$\overline{RA_w} = \frac{\sum_{r=1}^{23} w_{reg} \cdot RA_{reg}}{\sum_{r=1}^{23} w_{reg}}$$

To calculate mean rates of regional temperature and precipitation change, a least-squares fit was calculated for 20 years of initial condition ensemble mean yearly data. The rate of change is the slope of the best fit line and the weighted best fits for the rate of change values are calculated as with the other weighted means above.

Bibliography

The World Factbook. Washington, DC, 2009.

M.R. Allen and W.J. Ingram. Constraints on future changes in climate and the hydrologic cycle. *Nature*, 419:224–232, 2002.

Myles Allen. Do-it-yourself climate prediction. *Nature*, 401:642, 1999.

Timothy Andrews, Piers M. Forster, and Jonathan M. Gregory. A surface energy perspective on climate change. *Journal of Climate*, 22(10):2557–2570, 2009.

ASOC. Asilomar scientific organizing committee. the asilomar conference recommendations on principles for research into climate engineering techniques. November 2010.

G. Bala, P. B. Duffy, and K. E. Taylor. Impact of geoengineering schemes on the global hydrological cycle. *PNAS*, 105:7664–7669, 2008.

Richard Elliot Benedick. Considerations on governance for climate remediation technologies: Lessons from the "ozone hole". *STANFORD JOURNAL OF LAW, SCIENCE & POLICY*, IV:6–9, 2011.

J Eric Bickel and Lee Lane. An analysis of climate engineering as a response to climate change. Technical Report AR10-182, Copenhagen Consensus Center, Frederiksberg, Denmark, September 2009.

J. J. Blackstock, D. S. Battisti, K. Caldeira, D. M. Eardley, J. I. Katz, D. W. Keith, A. A. N. Patrinos, D. P. Schrag, R. H. Socolow, and S. E. Koonin. Climate engineering responses to climate emergencies. Technical Report archived online at: http://arxiv.org/pdf/0907.5140, Novim, Santa Barbara, California, July 2009.

O. Boucher, J.A. Lowe, and C.D. Jones. Implications of delayed actions in addressing carbon dioxide emission reduction in the context of geo-engineering. *Climatic Change*, 92: 261–273, 2009.

Leo Breiman, Jerome H. Freidman, Richard A. Olshen, and Charles J. Stone. *Classification and Regression Trees*. Chapman & Hall, New York, NY, 1984.

P. Brohan, J. J. Kennedy, I. Harris, S. F. B. Tett, , and P. D. Jones. Uncertainty estimates in regional and global observed temperature changes: A new data set from 1950. *J. Geophys. Res.*, 111, 2006.

Ken Caldeira and Lowell Wood. Global and arctic climate engineering: numerical model studies. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, 366(1882):4039–4056, November 2008.

M. Collins, S. F. B. Tett, and C. Cooper. The internal climate variability of hadcm3, a version of the hadley centre coupled model without flux adjustments. *Climate Dynamics*, 17(1): 61–81, 2001.

Paul J. Crutzen. Albedo enhancement by stratospheric sulfur injections: A contribution to resolve a policy dilemma? *Climatic Change*, 77(3-4):211–220, July 2006.

S.C. Doney, V.J. Fabry, R.A. Feely, and J.A. Kleypas. Ocean acidification: The other co2 problem,. *Annu. Rev. Marine. Sci.*, 1:169–192, 2009.

P. Forster and M. Collins. Quantifying the water vapour feedback associated with post-pinatubo cooling. *Climate Dynamics*, 23:207–214, 2004.

P. Forster, V. Ramaswamy, P. Artaxo, T. Berntsen, R. Betts, D.W. Fahey, J. Haywood, J. Lean, D.C. Lowe, G. Myhre, J. Nganga, R. Prinn, G. Raga, M. Schulz, and R. Van Dorland. *Changes in Atmospheric Constituents and in Radiative Forcing. In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2007. ISBN 0521880092.

D J Frame, T Aina, C M Christensen, N E Faull, S H E Knight, C Piani, S M Rosier, K Yamazaki, Y Yamazaki, and M R Allen. The climateprediction.net bbc climate change experiment: design of the coupled model ensemble. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, 367(1890):855–870, March 2009.

Stephen M. Gardiner. Is 'arming the future'with geoengineering really the lesser evil? some doubts about the ethics of intentionally manipulating the climate system. In Stephen M. Gardiner, Simon Caney, Dale Jamieson, and Henry Shue, editors, *Climate Ethics*. Oxford University Press, Oxford, UK, 2010.

Stephen M. Gardiner. Some early ethics of geoengineering the climate: A commentary on the values of the royal society report. *Environmental Values*, 20:163–188, 2011.

F. Giorgi and R. Francisco. Uncertainties in regional climate change prediction: a regional analysis of ensemble simulations with the hadcm2 coupled aogcm. *Climate Dynamics*, 16(2-3):169–182, 2000.

C. Gordon, C. Cooper, C. A. Senior, H. Banks, J. M. Gregory, T. C. Johns, J. F. B. Mitchell, and R. A. Wood. The simulation of sst, sea ice extents and ocean heat transports in a version of the hadley centre coupled model without flux adjustments. *Climate Dynamics*, 16(2-3): 147–168, 2000.

B. Govindasamy and K. Caldeira. Geoengineering earth's radiation balance to mitigate co2-induced climate change. *Geophysical Research Letters*, 27:2141–2144, 2000.

D Gregory and P R Rowntree. A mass flux convection scheme with representation of cloud ensemble characteristics and stability-dependent closure. *Monthly Weather Review*, 118: 1483–1506, 1990.

Lianhong Gu, Dennis D. Baldocchi, Steve C. Wofsy, J. William Munger, Joseph J. Michalsky, Shawn P. Urbanski, and Thomas A. Boden. Response of a deciduous forest to the mount pinatubo eruption: Enhanced photosynthesis. *Science*, 299(5615):2035–2038, 2003. doi: 10.1126/science.1078366. URL `http://www.sciencemag.org/content/299/5615/2035.abstract`.

J. Hansen, L. Nazarenko, R. Ruedy, Mki. Sato, J. Willis, A. Del Genio, D. Koch, A. Lacis, K. Lo, S. Menon, T. Novakov, Ju. Perlwitz, G. Russell, G.A. Schmidt, and N. Tausnev. Earth's energy imbalance: Confirmation and implications. *Science*, 308:1431–1435, 2005.

P. Heckendorn, D. Weisenstein, S. Fueglistaler, B. P. Luo, E. Rozanov, M. Schraner, L. W. Thomason, and T. Peter. The impact of geoengineering aerosols on stratospheric temperature and ozone. *Environ. Res. Lett.*, 4:045108, 2009.

O. Hoegh-Guldberg, P. J. Mumby, A. J. Hooten, R. S. Steneck, P. Greenfield, E. Gomez, C. D. Harvell, and et al. Coral reefs under rapid climate change and ocean acidification. *Science*, 318(5857):1737–1742, 12 2007.

L.W. Horowitz. Past, present, and future concentrations of tropospheric ozone and aerosols: Methodology, ozone evaluation, and sensitivity to aerosol wet removal. *J. Geophys. Res.*, 111:D22211, 2006.

Joshua B. Horton. Geoengineering and the myth of unilateralism: Pressures and prospects for international cooperation. *STANFORD JOURNAL OF LAW, SCIENCE & POLICY*, IV:56–69, 2011.

William Ingram. personal communication, 2008.

IPCC. *Climate Change 2007 - The Physical Science Basis: Working Group I Contribution to the Fourth Assessment Report of the IPCC*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2007a. ISBN 0521880092. URL `http://www.worldcat.org/isbn/0521880092`.

IPCC. *Climate Change 2007 - Impacts, Adaption and Vulnerability: Working Group II Contribution to the Fourth Assessment Report of the IPCC*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2007b.

Peter J. Irvine, Andy Ridgwell, and Daniel J. Lunt. Assessing the regional disparities in geoengineering impacts. *Geophysical Research Letters*, 37:L18702, 2010.

Micheline R. Ishay. *The History of Human Rights: From Ancient Times to the Globalization Era*. University of California Press, Berkeley and Los Angeles, California, 2004.

Dale Jamieson. Ethics and intentional climate change. *Climatic Change*, 77(3):323–336, 2006.

A. Jones, J. Haywood, O. Boucher, B. Kravitz, and A. Robock. Geoengineering by stratospheric so2 injection: Results from the met office hadgem2 climate model and comparison with the goddard institute for space studies modele. *Atmos. Chem. Phys. Discuss.*, 10:7421–7434, 2010.

David Keith. The case for geoengineering research. Presentation at MIT, Cambridge, MA, October 2009.

David W. Keith. Geoengineering the climate: History and prospect. *Annual Review of Energy and the Environment*, 25(1):245–284, 11 2000a.

David W. Keith. The earth is not yet an artifact. *IEEE Technology and Society Magazine*, 19: 25–28, 2000b.

David W. Keith. Photophoretic levitation of engineered aerosols for geoengineering. *PNAS*, 107(38):16428–16431, 2010.

David W. Keith, Edward Parson, and M. Granger Morgan. Research on global sun block needed now. *Nature*, 463(28):426–427, 2010.

Robert O. Keohane and David G. Victor. The regime complex for climate change. Technical Report Discussion Paper 2010-33, Harvard Project on International Climate Agreements, Cambridge, Mass., January 2010.

Daniel B. Kirk-Davidoff, Eric J. Hintsa, James G. Anderson, and David W. Keith. The effect of climate change on ozone depletion through changes in stratospheric water vapour. *Nature*, 402:399–401, 1999.

Christopher G. Knight, Sylvia H. E. Knight, Neil Massey, Tolu Aina, Carl Christensen, Dave J. Frame, Jamie A. Kettleborough, Andrew Martin, Stephen Pascoe, Ben Sanderson, David A. Stainforth, and M. R. Allen. Association of parameter, software, and hardware variation with large-scale behavior across 57,000 climate models. *Proceedings of the National Academy of Sciences*, 104(30):12259–12264, 2007.

Ben Kravitz, Alan Robock, Olivier Boucher, Hauke Schmidt, Karl E. Taylor, Georgiy Stenchikov, and Michael Schulz. The geoengineering model intercomparison project (geomip). *Atmospheric Science Letters*, 12(2):162–167, 2011. doi: 10.1002/asl.316.

C. Körner and F.A. Bazzaz. *Carbon dioxide, populations, and communities*. Academic Press, New York, New York, 1996.

F.H. Lambert, P.S. Stott, M.R. Allen, and M.A. Palmer. Detection and attribution of changes in 20th century land precipitation. *Geophys. Res. Lett.*, 31:L10203, 2004.

J. Latham, P. Rasch, C.C. Chen, L. Kettles, A. Gadian, A. Gettelman, H. Morrison, K. Bower, and T. Choularton. Global temperature stabilization via controlled albedo enhancement of low-level maritime clouds. *Phil. Trans. R. Soc. A*, 366:3969–3987, 2008.

T. M. Lenton and N. E. Vaughan. The radiative forcing potential of different climate geoengineering options. *Atmos. Chem. Phys.*, 9:5539–5561, 2009.

Timothy M. Lenton, Hermann Held, Elmar Kriegler, Jim W. Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhube. Tipping elements in the earth's climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793, 6 2008.

David B. Lobell, Wolfram Schlenker, and Justin Costa-Roberts. Climate trends and global crop production since 1980. *Science*, 2011. doi: 10.1126/science.1204531. Sciencexpress prepublication version.

D. J. Lunt, A. Ridgwell, P. J. Valdes, and A. Seale. 'sunshade world': A fully coupled gcm evaluation of the climatic impacts of geoengineering. *Geophys. Res. Lett.*, 35(12):L12710, 2008.

D. MacMynowski, H. Shin, K. Caldeira, and D. Keith. Can we test geoengineering? *Submitted*.

H. D. Matthews and K. Caldeira. Transient climate carbon simulations of planetary geoengineering. *Proceedings of the National Academy of Sciences*, 104(24):9949–9954, 6 2007.

Justin McClellan, James Sisco, Brandon Suarez, and Greg Keogh. Geoengineering cost analysis. Technical Report AR10-182, Aurora Flight Sciences Corporation, Cambridge, Mass., October 2010.

G.A. Meehl, T.F. Stocker, W.D. Collins, P. Friedlingstein, A.T. Gaye, J.M. Gregory, A. Kitoh, R. Knutti, J.M. Murphy, A. Noda, S.C.B. Raper, I.G. Watterson, A.J. Weaver, and Z.-C. Zhao. *Global Climate Projections. In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2007. ISBN 0521880092.

Juan Moreno-Cruz and David W. Keith. Climate policy under uncertainty a case for geoengineering. *Climatic Change*, under review.

Juan Moreno-Cruz, Katharine Ricke, and David W. Keith. A simple model to account for regional inequalities in the effectiveness of solar radiation management. *Climatic Change*, 2011. in press.

M. G. Morgan and David Keith. Improving the way we think about projecting future energy use and emissions of carbon dioxide. *Climatic Change*, 90:189–215, 2008.

M. G. Morgan and Katharine Ricke. Cooling the earth through solar radiation management: The need for research and an approach to its governance. Technical report, International Risk Governance Council, 2010.

M. Granger Morgan. Managing carbon from the bottom up. *Science*, 289:2285, September 2000.

M. Granger Morgan, Milind Kandlikar, James Risbey, and Hadi Dowlatabadi. Why conventional tools for policy analysis are often inadequate for problems of global change. *Climatic Change*, 41:271–281, 1999.

David R. Morrow, Robert E. Kopp, and Michael Oppenheimer. Toward ethical norms and institutions for climate engineering research. *Environ. Res. Lett.*, 4:045106, 2009.

Richard H. Moss, Jae A. Edmonds, Kathy A. Hibbard, Martin R. Manning, Steven K. Rose, Detlef P. van Vuuren, Timothy R. Carter, Seita Emori, Mikiko Kainuma, Tom Kram, Gerald A. Meehl, John F. B. Mitchell, Nebojsa Nakicenovic, Keywan Riahi, Steven J. Smith, Ronald J. Stouffer, Allison M. Thomson, John P. Weyant, and Thomas J. Wilbanks. The next generation of scenarios for climate change research and assessment. *Nature*, 463: 747–756, 2010.

J. M. Murphy, D. M. H Sexton, D. N. Barnett, G. S. Jones, M. J. Webb, M. Collins, and D. A. Stainforth. Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature*, 430(7001):768–772, 2004.

N. Nakiçenoviç and R.G. Swart. Cambridge, UK and New York, NY, USA.

NAS. 2101 Constitution Ave., N.W. • Washington, D.C. 20418, 1992. ISBN 0-309-04386-7.

William Nordhaus. Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences*, (103):3510–3517, 2006.

William Nordhaus. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, 2008.

Luke Oman, Alan Robock, Georgiy L. Stenchikov, and Thorvaldur Thordarson. High-latitude eruptions cast shadow over the african monsoon and the flow of the nile. *Geophys. Res. Lett.*, 33:L18711, 2006. doi: 10.1029/2009GL039209.

NRC Committee on Stabilization Targets for Atmospheric Greenhouse Gas Concentrations. *Climate Stabilization Targets: Emissions, concentrations, and Impacts over Decades to Millennia.* The National Academies Press. ISBN 0-30915177-5. URL `http://www.nap.edu/catalog/12877.html`.

J. R. Pierce, D. K. Weisenstein, P. Heckendorn, T. Peter, and D. W. Keith. Efficient formation of stratospheric aerosol for climate engineering by emission of condensible vapor from aircraft. *Geophys. Res. Lett.*, 37:L18805, 2010.

V. Ramaswamy, O. Boucher, J. Haigh, D. Hauglustaine, J. Haywood, G. Myhre, T. Nakajima, G.Y. Shi, and S. Solomon. *Radiative Forcing of Climate Change In: Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2001.

D.A. Randall, R.A. Wood, S. Bony, R. Colman, T. Fichefet, J. Fyfe, V. Kattsov, A. Pitman, J. Shukla, J. Srinivasan, R.J. Stouffer, A. Sumi, and K.E. Taylor. *Climate Models and Their Evaluation. In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2007. ISBN 0521880092.

N. A. Rayner, P. Brohan, D. E. Parker, C. K. Folland, J. J. Kennedy, M. Vanicek, T. J. Ansell, and S. F. B. Tett. Improved analyses of changes and uncertainties in the sea surface temperature measured in situ since the mid-nineteenth century: The hadsst2 dataset. *J. Clim.*, 19:446–469, 2006.

S. Rayner, Redgwell C., J. Savulescu, N. Pidgeon, and T. Kruger. Memorandum on draft principles for the conduct of geoengineering research. House of Commons Science and Technology Committee enquiry into The Regulation of Geoengineering., 2009.

Kate Ricke, Dan Rowlands, William Ingram, David Keith, and M. Granger Morgan. How does the sensitivity of climate affect stratospheric solar radiation management? *in preparation for Nature Climate Change*, 2011.

Katharine Ricke, Granger Morgan, and Myles Allen. Regional climate response to solar radiation management. *Nature Geoscience*, 3(8):537–541, 2010.

A. Robock, A. Marquardt, B. Kravitz, and G. Stenchikov. Benefits, risks, and costs of stratospheric geoengineering. *Geophys. Res. Lett.*, 36:L19703, 2009. doi: 10.1029/2009GL039209.

Alan Robock. Volcanoes: Role in climate. In J. Holton, J. A. Curry, and J. Pyle, editors, *Encyclopedia of Atmospheric Sciences*, pages 2494–2500. Academic Press, London, 2003.

Alan Robock, Luke Oman, and Georgiy L. Stenchikov. Regional climate responses to geoengineering with tropical and arctic so2 injections. *Journal of Geophysical Research*, 113 (D16), 8 2008.

Alan Robock, Martin Bunzl, Ben Kravitz, and Georgiy L. Stenchikov3. A test for geoengineering? *Science*, 327(5965):530–531, 2010. doi: 10.1126/science.1186237.

Gerard H. Roe and Marcia B. Baker. Why is climate sensitivity so unpredictable? *Science*, 318:629–632, 2007.

D. Rowlands, D. Frame, M. Allen, T. Aina, and M. Thurston. Quantification of uncertainties in 21st centure temperature projections. *In Draft*, 2011.

Ian H. Rowlands. Classical theories of international relations. In Urs Luterbacher and Detlef Sprinz, editors, *International Relations and Global Climate Change*, pages 43–65. 2001.

S. Salter, G. Sortino, and J. Latham. Sea-going hardware for the cloud albedo method of reversing global warming. *Phil. Trans. R. Soc. A*, 366:3989–4006, 2008. doi: 10.1098/rsta.2008.0136.

M. Sato, J.E. Hansen, M.P. McCormick, and J.B. Pollack. Stratospheric aerosol optical depth, 1850-1990. *J. Geophys. Res.*, 98:22987–22994, 1993.

John Shepherd and Steven Rayner. Responses to climate change: the four-fold way. submitted.

John Shepherd, Ken Caldeira, Joanna Haigh, David Keith, Brian Launder, Georgiina Mace, Gordon MacKerron, John Pyle, Steve Rayner, Catherine Redgwell, and Andrew Watson. Geoengineering the climate: science, governance and uncertainty. Technical report, Royal Society, 2009. URL ``http://royalsociety.org/document.asp?tip=0&id=8770''.

Jack Snyder. One world, rival theories. *Foreign Policy*, 145:52–62, 2004.

B.J. Soden, Richard T. Wetherald, Georgiy L. Stenchikov, and Alan Robock. Global cooling after the eruption of mount pinatubo: A test of climate feedback by water vapor. *Science*, 296:727–730, 2002.

S. K. Solanki and N. A. Krivova. Can solar variability explain global warming since 1970? *J. Geophys. Res.*, 108(A5):1200, 2003.

Detlef Sprinz and Tapani Vaahtoranta. The interest-based explanation of international environmental policy. *International Organization*, 48:77–105, 1994.

Detlef F. Sprinz. Comparing the global climate regime with other global environmental accords. In Urs Luterbacher and Detlef Sprinz, editors, *International Relations and Global Climate Change*, pages 247–277. 2001.

D. A. Stainforth, T. Aina, C. Christensen, M. Collins, N. Faull, D. J. Frame, J. A. Kettleborough, S. Knight, A. Martin, J. M. Murphy, C. Piani, D. M. H Sexton, L. A. Smith, R. A Splcer, A. J. Thorpe, and M. R. Allen. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, 433(7024):403–406, 2005.

David Stainforth, Jamie Kettleborough, Andrew Martin, Andrew Simpson, Richard Gillis, Ali Akkas, Richard Gault, Mat Collins, David Gavaghan, and M. R. Allen. Climateprediction.net: design principles for public resource modelling research. In *Proc. 14th IASTED conference on parallel and distributed computing systems*, 2002.

G. Stenchikov, K. Hamilton, R. J. Stouffer, A. Robock, V. Ramaswamy, B. Santer, and H.-F. Graf. Arctic oscillation response to volcanic eruptions in the ipcc ar4 climate models. *J. Geophys. Res.*, 111:D07107, 2006.

E. Teller, L. Wood, and R. Hyde. Global warming and ice ages: Prospects for physics based modulation of global change. Technical Report UCRL-JC-128715, Lawrence Livermore National Laboratory, August 1997.

E. Teller, R. Hyde, and L. Wood. Active climate stabilization: Practical physics-based approaches to prevention of climate change. Technical Report UCRL-JC-148012, Lawrence Livermore National Laboratory, April 2002.

Simone Tilmes, Rolando R. Garcia, Douglas E. Kinnison, Andrew Gettelman, and Philip J. Rasch. Impact of geoengineered aerosols on the troposphere and stratosphere. *Journal of Geophysical Research*, 114(D12), 6 2009.

K. E. Trenberth and A. Dai. Effects of mount pinatubo volcanic eruption on the hydrological cycle as an analog of geoengineering. *Geophys. Res. Lett.*, 34:L15702, 2007.

S. Twomey. The influence of pollution on the shortwave albedo of clouds. *Journal of the atmospheric sciences*, 34:1149–1153, 1977.

D Victor, MG Morgan, J Apt, J Steinbruner, and K Ricke. The geoengineering option. *Foreign Affairs*, 88:64–76, 2009.

David G. Victor. *Global Warming Gridlock*. Cambridge University Press, 2011.

Paul Viotti and Mark Kauppi. *International Relations Theory*. Longman, New York, NY, USA, 2010. ISBN 0131892614.

Martin L. Weitzman. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, XCI(1), 2009.

Nicholas J. Wheeler. *Saving strangers: humanitarian intervention in international society*. Oxford University Press, Oxford, UK, 2000.

Ellen Wiegandt. Climate change, equity, and international negotiations. In Urs Luterbacher and Detlef Sprinz, editors, *International Relations and Global Climate Change*, pages 127–150. 2001.

T. M. L. Wigley. A combined mitigation/geoengineering approach to climate stabilization. *Science*, 314(5798):452–454, 10 2006.

Hiro Yamazaki. personal communication, 2008.

Fanglin Yang, Arun Kumar, Michael E. Schlesinger, and Wanqiu Wang. Intensity of hydrological cycles in warmer climates. *Journal of Climate*, 16(14):2419–2423, 2003.