The Report Committee for Michelle Carolina Buckholtz
Certifies that this is the approved version of the following report:

Decision Analysis

for Climate Engineering Research

APPROVED BY

SUPERVISING COMMITTEE:

Supervisor: ______________________________
Sheila Olmstead

Co-Supervisor: ______________________________
J. Eric Bickel

Decision Analysis

for Climate Engineering Research

by

Michelle Carolina Buckholtz, B.A.

Report

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Master of Public Affairs

The University of Texas at Austin

May 2015

Abstract

Decision Analysis

for Climate Engineering Research

Michelle Carolina Buckholtz, MPAff

The University of Texas at Austin, 2015

Supervisors: Sheila Olmstead, J. Eric Bickel

Technology solutions designed to manage climate change risk fall into three categories: mitigation, adaptation, and climate engineering. While mitigation and adaptation technologies are well established and have substantial public support as policy alternatives, climate engineering strategies remain mostly in the early stages of research and development. Both the further pursuit of research and eventual use of climate engineering technologies have been subject to moral and ethical objections. The intention of this report is to aid policy-makers in the decision as to whether society should pursue climate engineering research. This report identifies the unique characteristics which make climate engineering an important tool in the portfolio of strategies for managing climate change risks. Next potential benefits and costs associated with the technology are explored. The largest ethical objections to research and use of the technology are discussed and presented in a more consistent framework than found in existing literature. Finally, a model evaluating the sensitivity of the decision to pursue

climate engineering research to two large ethical objections was built. Using outputs from an existing climate model, the analysis in this report adjusts the likelihood of the two ethical objections occurring across several scenarios to illustrate how the quality of the decision changes based on different assumptions about society.

Table of Contents

Chapter 1: Introduction

Current climate change policies are not on track to limit global temperature change to the 2° Celsius goal set by the Intergovernmental Panel on Climate Change (IPCC). It has been estimated that "a continuation of business as usual trends will likely lead to a [tripling] of CO_2 by the end of the century, with a 50% chance of exceeding the rise of five degrees Centigrade, about the same as the increase from the last ice age to the present" [1]. There is an increasing demand for innovation and technological solutions to fix this problem and lower the risks associated with climate change and the rise in global temperatures. Researchers have proposed climate engineering technologies as part of the portfolio of tools to manage climate change risks.

The three major approaches to the challenge of managing climate change risks are mitigation, adaptation, and climate engineering. Climate change mitigation and adaptation technologies are considerably more developed than climate engineering technologies, and they are also more widely acknowledged as acceptable approaches to combat climate change risks. However, this has not always been the case. As recently as the 1990s, adaptation strategies were regarded as unacceptable and characterized by laziness [2]. Still, in its most recent publications the IPCC includes only mitigation and adaptation technologies in its set of alternatives for managing climate change risks [3].

Adaptation can be described as "a series of measures or actions to reduce the negative impacts of climate change… and to take advantage of opportunities it may offer" [4]. This includes both reactive and anticipatory measures in all imaginable fields. Brasseur and Granier [4] consider a diverse portfolio of adaptation measures ranging from protection against sea level rise in coastal areas to healthcare systems which can adequately manage changing disease patterns.

Climate change mitigation technologies include any that address reducing carbon dioxide emissions. This large category of mitigation technologies has been split into two sub-categories: energy production and energy use and consumer behavior. Energy production technologies include solar energy, wind energy, hydropower, nuclear energy, and geothermal energy [5]. Arvesen, Bright, and Hertwich [6] challenge the "unfounded technology optimism" in these clean energy technologies. The authors argue that transition to clean energy will in itself cause climate impacts and that developing fossil fuels in tandem with renewable energy may lower system performance and lead to carbon lock-in. Further, consumer behavior and energy efficiency improvements are likely less effective than predicted due to possible rebound effects [6].

Brasseur and Granier [4] estimate that a decrease of roughly 80% of current emissions is required to remain under the target 2°C temperature increase. If emissions reductions were on target by 2011, this would equate to a roughly 3.7% decrease per year, but this value increases to nearly 10% per year if business-as-usual trends continue until 2020 [4]. Although many countries support international targets for limiting temperature change, emissions reductions are especially vulnerable to free riding [7]. Without proper incentives to encourage compliance with voluntary international agreements, the demands of such agreements will not be met. Four years after the 3.7% annual emissions reductions should have begun, emissions reductions are still increasing to accommodate the energy demands of large developing countries like China and India. A successful mitigation strategy would also need to include practical ways for China, India, and other emerging economies to reduce emissions cheaply as huge cost penalties are incurred when there is less than universal participation in a mitigation program [8].

The umbrella term 'mitigation and adaptation technologies' describes a very broad set of technologies with loose boundaries. To better understand how a technology would be categorized as a climate engineering technology, this large category is contrasted with mitigation and adaptation technologies. Several texts include carbon dioxide removal technologies in the mitigation category, and this is also a fair way to conceptualize the technologies. Keith [9] illustrates this point when he notes, "until humanity's net emissions are zero, any carbon removal method has precisely the same effect on the climate as mitigation – a ton not emitted is the same as a ton emitted and recaptured." In this report, mitigation has the somewhat more rigid definition as any technology concerning itself with emissions reductions. Climate engineering technologies considered in this report are sorted into two technology categories and shown in Table 1.[1]

Solar Radiation Management (SRM)	Carbon Dioxide Removal (CDR)
Sulfate aerosols Stratospheric aerosols Tropospheric aerosols **Sunshades in space** **Enhanced cloud + surface albedo** Mechanically enhanced cloud albedo Biologically enhanced cloud albedo Increasing grassland and cropland albedo Increasing human settlement albedo Increasing desert albedo	**Air capture + storage** **Land carbon sink enhancement** Afforestation and reforestation Bio-char Bio-energy with carbon capture + storage **Ocean carbon sink enhancement** Enhancing the solubility pump Increasing ocean alkalinity Enhancing the biological pump Iron fertilization Macronutrient fertilization Enhancing upwelling

Table 1: Proposed Climate Engineering Approaches by Technology Category [4]

[1] Table 1 is a modification to a table provided by Guy Brasseur and Clair Granier in *Mitigation, Adaptation or Climate Engineering*.

There are many ways to separate technologies into an appropriate approach category, and each has its potential merits and biases. The key factor for defining categories in this report is intentional manipulation of the climate system. Certainly emissions reduction technologies also exist to lessen the impacts on the climate system. Under this categorization method however, that goal is considered secondary and the technology is placed in the mitigation category. The intention of categorizing the technologies in this manner is twofold. It provides an interesting way to think about the categories as they develop, and it loosely follows the categorization of previous studies relevant to this paper.

The intention of this report is to aid policy-makers in the decision as to whether society should pursue climate engineering research. This report considers the role of climate engineering technologies in a portfolio for managing climate change risks through a comparative analysis of current and proposed climate engineering technologies in Chapter Two. Once the focus has been narrowed to a single technology, sulfate aerosols, potential benefits and costs of researching and using the technology are identified in Chapter Three. Numerous studies provide ethical and moral critiques of the proposed technology, often including cost-benefit estimations. Chapter Four identifies the largest ethical and moral objections to pursuing climate engineering research and attempts to frame each in a manner consistent with society's objectives for managing climate change risks.

Two of the largest objections to researching the technology, that climate engineering research would lessen society's mitigation efforts and would necessarily lead to deployment of the technology, are investigated further in Chapter 5. Using outputs from existing cost-benefit models for climate change, the sensitivity of the decision is tested by adjusting the likelihood of each of these two objections. Finally, Chapter Six synthesizes all of the previous work to relate ideas back to the decision problem at hand.

Chapter 2: Introduction to Climate Engineering

Looking back not even twenty years, climate engineering was such an unfamiliar concept that researchers had difficulty describing what climate engineering encompassed. Thomas Schelling [10] described it as "a new term, still seeking a definition. It seems to imply something global, intentional, and unnatural." Over the years a consensus appears to have developed with climate engineering frequently defined as "the deliberate, large-scale manipulation of the planetary environment to counteract anthropogenic climate change" [4, 11]. While this broad category is applicable to many technologies, it is easily separated into two large technology sub-categories: Solar Radiation Management (SRM) technologies and Carbon Dioxide Removal (CDR) technologies.

This chapter takes a closer look at each of the technologies in Table 1, exploring how the technologies are expected to work, the timescale for their effectiveness, and roughly how much is understood about the technologies.

2.1 Solar Radiation Management

As seen in Table 1 SRM technologies include sulfate aerosols, sunshades in space, and enhanced cloud and surface albedo. Though different in implementation, each of these technologies work by reflecting some amount of solar radiation back to space, effectively cooling the Earth [12]. Studies estimate that a 1.8% decrease of all incoming solar radiation would offset a doubled pre-industrial atmospheric concentration of carbon dioxide [13]. Temperatures could then return to pre-industrial levels [14].

2.1.1 Sulfate Aerosols

Sulfate aerosols work to cool the Earth by two means: they directly scatter a portion of solar radiation back to space and they indirectly alter the optical and physical properties of clouds [4]. A common example used to explain sulfate aerosols is a volcanic eruption. When a volcano erupts SO_2 is released and is converted into tiny sulfate particles in the stratosphere. These particles reflect the incoming sunlight and cool the Earth by blocking out a small portion of the incoming heat from the sun. The effects of volcanic eruptions have been studied, and an average surface cooling of 0.5°C was observed in the year following the large eruption of Mount Pinatubo in 1991 [14, 15]. The idea behind sulfate aerosol injections is to simulate natural and human processes which currently release sulfates into the air – as many fossil-fuel power plants and volcanic eruptions do – but with measured amounts of sulfate particles released at given time frames to maintain a cooler temperature as the Earth warms.

The timescale for sulfate aerosols to work is relatively quick. As with the volcano example, measurements of the Earth's average surface temperature were 0.5°C less within just one year. The issue with time for sulfate aerosols is more of a matter of duration than onset of effectiveness. In the troposphere the particles remain for roughly a few days to one week, and in the stratosphere, a further layer of the atmosphere, the particles remain between 1 and 2 years [9, 15]. Researchers propose that this technology has high potential to be technically possible; however, most of the research on sulfate aerosols exists only on paper [16, 17].

2.1.2 Sunshades in Space

Sunshades in space have also been characterized as "mirrors" in space [14, 18]. Sunshades could be placed in an orbit around the sun or around the Earth, with the magnitude of the negative radiative forcing effect dependent on the size and placement of the shades [13]. These sunshades would function somewhat similarly to the sulfate particles, but rather than scattering sunlight with small reflective pieces of sulfate dust, it is more similar to placing a tarp or tent outdoors for a larger block of shade in one concentrated area.

As with the aerosols, sunshades have a relatively quick onset for effectiveness. The problem arises when consideration for non-static radiative forcing is given. More simply, that means that because greenhouse gas emissions continue, the amount of warming that needs to be countered is also increasing. The amount of shade needed to offset the warming then also increases over time. In order to remain effective, an estimated few hundred thousand launches into space each year would be needed to increase the amount of sunshades in orbit [13].

2.1.3 Enhanced Cloud + Surface Albedo

Albedo refers to the ratio of light reflected away from the Earth to the amount received from the Sun. Technologies for enhancing cloud and surface albedo work then to reflect sunlight and the heat from solar radiation back to space. Stephen Salter has proposed unmanned vessels which could travel across the oceans powered by wind. The vessels would spray filtered water from the ocean into existing clouds to brighten them and thereby increase the amount of solar radiation reflected back to space [19]. This method is expected to be very inexpensive. The amount of vessels needed would depend on the desired temperature as well as the rate of emissions and mitigation, but estimates mark 1,500 vessels as a starting point with an additional 50 each year to account for continued emissions [14, 19].

2.2 Carbon Dioxide Removal

As seen in Table 1 CDR technologies include air capture and storage, land carbon sink enhancement, and ocean carbon sink enhancement. These technologies function essentially as their name leads one to imagine; they remove carbon dioxide from the air.

2.2.1 Air Capture + Storage

Currently, the most effective methods of CDR involve capturing the CO_2 from the most concentrated sources, placing it under a large amount of pressure to turn the gas into a liquid[2], and then injecting it deep underground for long-term storage [20]. This is the underlying concept of air capture and storage technologies, sometimes called carbon capture and storage (CCS) technologies. Fossil fuel power plants are key points of interest for CCS because of the concentrated amounts of CO_2 present in comparison with ambient air. Other CCS research looks to more location-independent means of capturing carbon dioxide. Though independent of existing energy infrastructure systems, these technologies function very similarly [16].

CCS technologies are the most developed of the climate engineering technologies discussed so far, but they are arguably still in an early stage of development. In a study looking at the invention and transfer rates of climate change mitigation technologies, it was found that "the average export rate of CCS inventions[3] was 20.5% from 2000 to 2006, significantly above the rate for other climate-mitigation technologies (15 percent), suggesting a higher quality of patented inventions, which is consistent with an early stage of technology development" [21].

While these inventions are developed and transferred at an increasing rate, the time until they are effective is hindered in two ways. The first obstacle is a technological concern regarding site selection for injecting the CO_2 [20]. Detailed exploration could resolve this obstacle in a relatively short time, perhaps on the order of one year to one decade. The second obstacle occurs due to the increasing rate of greenhouse gas

[2] For the technical reader, this is actually a 'supercritical fluid' which has properties similar to both liquids and gases.

[3] The average export rate of CCS inventions refers to the share of inventions which have patents filed in multiple locations. For example, if one invention had a patent filed in both Germany and China and three inventions had patents filed only in Japan, the export rate would be 25%.

emissions offsetting the effects of captured carbon dioxide. Carbon dioxide has a long residence in the atmosphere, ranging from centuries to millennia [22]. Due to this long lifetime, the effectiveness of CCS technologies is determined by their scale relative to emissions. That is, taken as an independent tool CCS technologies would need to grow at least as quickly as the growth in emissions in order to be effective.

A final consideration about CCS technologies is the finite storage space available for the pressurized CO_2. Current estimates range from at least 100 years up to more than 300 years of storage space available [5, 20]. While these estimates are all contingent on the amount of carbon dioxide captured, policies and safety regulations in place, and the results of future exploration and technological advances, there is an eventual upper limit.

2.2.2 Land Carbon Sink Enhancement

Land carbon sink enhancement methods and technologies aim to increase the amount of carbon stored in soil and vegetation. Though the majority of carbon on land is stored within the soil, attention is often focused on afforestation and reforestation [13]. This is essentially a process for converting non-forested land to forested land, or more simply, planting trees. Studies have estimated the possible change in radiative forcing from successful afforestation, and while it appears to be an effective tool, the researchers expressed doubt with the reality of achieving those results [13].

Afforestation presents a number of potential adverse effects as well, including loss of species diversity from changing habitats and lower land surface albedo. Perhaps the largest concern regarding afforestation methods is fire. Forest fires are naturally occurring and necessary tools to maintaining a healthy forest. When the trees burn, all the stored carbon is released back into the atmosphere as part of its cycle.

Other methods of land carbon sink enhancement such as bio-char and bio-energy with carbon capture and storage use the biomass generated in the afforestation process. The biomass is burned in controlled manner, and the carbon is captured to prevent its release into the atmosphere [13]. In any of these scenarios, significant land use changes are required and act as the first key barrier to the effectiveness of these methods.

2.2.3 Ocean Carbon Sink Enhancement

Methods to enhance the ocean carbon sink include enhancing both the solubility pump and the biological pump. The solubility pump refers to the absorption of CO_2 by surface ocean waters as they cool moving from low to high latitudes. Potential enhancements to the solubility pump involve increasing the amount of CO_2 that is absorbed into the ocean and the rate at which it sinks into deeper water. To increase the rate of sinking, floating barges would create thicker sea ice and cool the water. This method has very high costs and uncertain effectiveness, leading some to consider this approach "wholly ineffectual." Proposals to increase the quantity absorbed focus on lowering the pH of the ocean by adding ground limestone to the water. When strategically dispersed, the increased absorption of CO_2 provided by interactions with the limestone shows potential for lowering atmospheric concentrations of CO_2 [13].

A more promising method of ocean carbon sink enhancement is iron fertilization, a method of augmenting the biological pump. The biological pump refers to carbon sinking into the lower ocean trapped in living and dead organisms. Iron fertilization adds this essential nutrient to ocean waters where other key nutrients already exist in an attempt to increase the amount of photosynthesis in living organisms at the surface. Atmospheric carbon is taken in and trapped by the process. Several iron fertilization experiments have been conducted, but the results are not entirely clear. Roughly half of the experiments showed an increase in the amount of carbon absorbed, but some argue this observation is due to inconsistent ocean conditions during evaluation periods [13].

Each of these ocean technologies face similar barriers. Storing carbon in the ocean is only a temporary solution, and the carbon will eventually return to the atmosphere. Additional concerns relate to the effects of ocean acidification and the impacts of these methods on marine species.

2.3 No Substitute Goods for SRM

Economists have established the role of climate engineering as a substitute for emissions reductions as both can counteract climate change [12]. This relationship is seemingly intuitive when looking through the lens of limiting climate change. However, the way this dynamic changes when considering certain climate engineering technologies not only as a tool to counteract climate change but also as insurance against damages from abrupt climate changes in a high climate sensitivity situation appears to be understated in much of the discussion. Brasseur and Granier [4] highlight the properties that give stratospheric aerosols an advantage over other climate engineering in this role by ranking each technology in effectiveness, affordability, timeliness, and safety. Table 2 below is an adapted version of the rankings provided in their paper [4].

Climate Change Technology Rankings for Effectiveness, Affordability, Timeliness, and Safety (1 = Lowest Ranking, 5 = Highest Ranking)				
Technology	**Effectiveness**	**Affordability**	**Timeliness**	**Safety**
Stratospheric aerosols	4	4	4	2
Sunshades in space	3	1.5	1	3
Enhanced cloud albedo	2.5	3	3	2
Surface albedo (urban)	1	1	3	5
Surface albedo (desert)	2.5	1	4	1
Afforestation and reforestation	2	5	3	4
Bio-energy with carbon capture + storage	2.5	2.5	3	4
Bio-char	2	2	2	3
Iron fertilization	2	3	1.5	1
Air capture + storage	4	1.9	2	5
CCS at source	3	3	4	5

Table 2: Climate Change Technology Rankings for Effectiveness, Affordability, Timeliness, and Safety [4]

In each column of the table scores range from 1 to 5, with 1 as the worst and 5 as the best ranking possible. Stratospheric aerosols stand out in all categories except for safety. The only technology included in the comparison with competitive rankings to aerosols is CCS at the source. While CCS technologies are well-established in comparison and incorporated into existing climate change governance regimes, they are unable to respond quickly to abrupt temperature increases. The changes in temperature over the next few decades are expected to be nearly independent of any changes in CO_2 concentrations mostly due to climate inertia [4, 9]. The ability to cheaply, effectively cool the surface temperature in a very short amount of time makes stratospheric aerosols unparalleled in the risk management realm. Arguably there currently exist no substitute goods which provide this insurance against abrupt climate change. This uniqueness gives reason to have interest in the technology – at least interest enough to pursue further analysis. For that reason, the remainder of this paper focuses on stratospheric aerosols. Hereafter the terms ‘SRM’ and ‘climate engineering’ will be used interchangeably to refer to stratospheric aerosols, both for simplification and to mirror the semantics of previous studies.

Chapter 3: Cost-Benefit Analysis for Climate Engineering Research

Climate engineering may prove itself a viable tool in our portfolio for managing climate change damages. However, the lack of a substitute good is not reason enough to pursue climate engineering research. This chapter introduces the main question of this paper: should society pursue climate engineering research.

Cost-benefit analysis is both a practical and commonly used tool to inform policy and decision making. Many researchers and decision makers, including the IPCC, acknowledge the limitations of cost-benefit analysis but are confident in its application to informing climate policy [3]. Still others are quite opposed. Stephen Gardener argues that the underlying economic theory for cost-benefit analysis is not developed enough for models to shape climate change policy. He quotes a moral philosophy professor on the topic as stating, "Cost-benefit analysis, when faced with uncertainties as big as these, *would simply be self-deception*" [23].

Despite this disagreement, it is still useful to identify and understand costs and benefits of potential SRM use. The following sections describe the costs and benefits of potential SRM use, with information about uncertainty and estimates from existing models. Some of the costs are categorically similar to benefits. Rather than describe each category as a benefit or cost based on net totals of the effects, some categories will appear in both sections to give a qualitative description of the desired or undesired effects.

3. 1 Benefits of Potential SRM Use

3.1.1 SRM is a Risk Management Tool in Tipping Point Scenarios

There is a large amount of uncertainty about climate sensitivity arising from factors such as the change in temperature due to a doubling of the amount of CO_2 in the atmosphere and the amount of time it takes CO_2 to leave the atmosphere among others [8]. These uncertainties make it difficult for a policy to assign a true upper limit of temperature change given certain actions. Bickel [18] simulated the probabilities of exceeding particular temperatures under different emissions mitigation policies. Even under a strict emissions control policy limiting temperature change to 1.5 °C there is still a nearly 20% chance of exceeding a temperature change of 2 degrees. More lenient policies had similar results; for example, the business as usual case had an 87% chance of temperature changes larger than 3 degrees with a maximum reaching temperature changes up to 10 °C [18].

Due to this great amount of uncertainty, researchers have proposed climate engineering as a tool for managing climate risk [9]. Rather than switching to a stricter mitigation policy with lower probabilities of exceeding high temperature changes, climate engineering allows the decision maker to remain in the same mitigation plan while removing the likelihood for passing certain temperatures. SRM is very fast acting, and estimates expect it to take about half a year to work effectively [15]. It serves as a quick fix or "to guarantee that the worst risks won't happen" [12].

3.1.2 Climate Engineering as a "Great Experiment"

Climate engineering has been described as a "great experiment" because many of the associated risks and benefits resemble that of a large-scale experiment. The main benefit from approaching SRM as an experiment is the limited risk, which is bounded by both the ability to stop the use and the short lifetime of aerosols in the atmosphere [12]. The effects of different methods are also able to be controlled based on the location of deployment. When used at higher latitudes, smaller areas are exposed to the experiment.

3.1.3 Improvements to Human Health

Climate engineering has been characterized as an “ugly technical fix,” where society intentionally pollutes to counter the effects of another pollutant [9]. Many would hesitate to attribute increased pollution with improvements to human health, but studies have found possible health benefits from SRM use. Perhaps the greatest benefit is that to agriculture. Allowing higher concentrations of CO_2 to remain in the atmosphere increases agricultural productivity [4, 12]. In addition to more productive crops, SRM may reduce the monsoon related risks to Asian agriculture by slowing or reducing the increases in precipitation predicted in climate change models [9]. Both of these factors contribute to a reliable food supply, a challenge to many poor and growing regions.

SRM also provides a more direct benefit to human health in its very basic function of reducing the amount of UV radiation reaching the Earth [12]. Similar human health benefits were observed with the 1987 Montreal Protocol, an international treaty which addressed ozone damaging emissions. Studies have investigated the extent of the benefits of this incredibly successful policy and identified substantial human health benefits in three categories: cataract incidence, cutaneous malignant melanoma (CMM) incidence and mortality, and non-melanoma skin cancer (NMSC) incidence and mortality. Current estimates show that “when compared with a situation of no policy controls, full implementation of the Montreal Protocol, including its Amendments and Adjustments, is expected to avoid more than 280 million cases of skin cancer, approximately 1.6 million skin cancer deaths, and more than 45 million cases of cataract in the United States for cohort groups in birth years 1890-2100” [24]. While benefits of this nature would differ in magnitude dependent on the amounts of SRM used and the duration of use, benefits in these same health categories could be expected.

3.2 Costs of Potential SRM Use

3.2.1 R&D and Implementation

There are multiple estimates for how much a climate engineering R&D program would cost, and most agree that SRM is an incredibly cheap alternative. The cost-benefit ratio for a climate engineering R&D program is on the order of 1000:1, and costs approximately $1 billion [9, 25].

3.2.2 Harms to Human and Ecological Health

As mentioned in the previous section on human health, climate engineering has been characterized as polluting to counter the effects of a different pollutant. Air pollution caused by the sulfate particles has adverse effects on human health. Estimates for premature deaths caused by current levels of outdoor air pollution and particulate matter range from roughly a half million to 1.3 million people annually [9, 15]. Effects from SRM use could potentially differ from these because the aerosols are injected into a higher section of the atmosphere and distributed in a more global than local manner [9]. Though estimates imply that net loss of human life would decrease, it is important to note that some amount of deaths could be directly attributable to increased air pollution from SRM use.

Sulfate air pollution has been attributed to ecological damages from acid rain and deposits [15]. SRM use could also contribute to ozone depletion, but this is not a necessary effect of the technology itself [15, 26]. Damages to the protective layer of ozone are caused by interactions of the sulfate particles with existing chlorine pollutants in the air. However, international treaty has successfully limited the use of the chlorine pollutant, and levels of chlorine remaining in the atmosphere continue to decline, making the effect of SRM use on ozone dependent on the timing of its use [9].

3.2.3 Changing Weather Patterns

Many concerns about climate engineering research arise due to the uncertain effects to temperature, precipitation, and climate patterns. Climate models include possible changes to large weather patterns like the El Niño and Asian and African summer monsoon [26]. Keith [9] argues however, that changes to these systems may not necessarily imply increased damages. Depending on the way in which SRM is used, changes to the monsoons could provide benefits to agricultural production in the region and decreased risk of major flood damages. The weather damages are difficult to predict and difficult to assign as a net cost or net benefit without understanding the extent and manner of SRM use.

Chapter 4: Ethical Considerations

This chapter presents the most prevalent ethical objections to climate engineering, and each is discussed as the arguments appear in the literature as well as how they are represented within the frame of our decision. There are two distinct categories of ethical and moral claims against climate engineering research: objections to climate engineering technologies themselves (the "what") and objections to governance, usage, and control of the technologies (the "how"). Most of the objections to SRM technologies arise due to the great uncertainties about climate sensitivity, society's preferences, and the extent of the effects of SRM technologies. Though distinct, many of these objections are presented together under a single argument. Where relevant, the categories of the objection will be noted in the following discussions of each objection.

4.1 Moral Hazard and Risk Compensation

Moral hazard has been described as "a socially inefficient increase in risk-taking by one party once another party absorbs some of the potential negative consequences of the first party's actions" [2]. This effect is often observed in situations where insurance is involved. For example, a driver might drive less cautiously if he knows the auto insurance company will pay for a large amount of possible damages. When considering climate engineering, moral hazard is also often presented as risk compensation or mitigation obstruction. This difference is due to the fact that the risk is not transferred to another party, but rather reduced by technological innovation [2, 27]. Though this use of the term moral hazard is not consistent with its typical association with a market failure, the terms will be treated as interchangeable in this report to be consistent with previous papers.

The moral hazard objection to pursuing climate engineering research identifies a trade-off between climate engineering and mitigation efforts. If SRM research happens, society will have less motivation to pursue costly mitigation efforts and will not control emissions as much as needed [2, 9, 12, 27-29]. Schelling [10] stated this idea very simply by noting that "some of us may not take global warming quite so seriously if we believe it may eventually be susceptible to direct intervention." Some arguments have called for a complete ban on SRM research so that society will only have the alternative to reduce emissions [12]. Others argue for precautions in pursuing SRM research rather than a complete research moratorium [27]. In our decision about pursuing climate engineering research, it is important to consider a few aspects of the moral hazard objection including the objectives in pursuing SRM research and emissions mitigation and the empirical evidence for moral hazard.

Society's objectives are very important for generating alternatives for addressing climate change. The IPCC has explicitly stated that the ultimate objective of their work is "the stabilization of greenhouse gas concentrations in the atmosphere…" [3]. Should society agree with this objective, there is little reason to pursue climate engineering research. Others however, have expressed an objective of reducing climate risks and damages or ensuring social welfare [2]. Because the stabilization of GHG concentrations in the atmosphere is one means of reducing climate risks and damages and ensuring social welfare, this analysis considers society's objective to be the latter.

With a clear objective of reducing climate risks in mind, the next appropriate question is whether a reduction in mitigation caused by the pursuit of SRM research necessarily decreases social welfare. David Morrow crafted an in-depth philosophical analysis of this question under various interpretations of moral thought including simple utilitarian, fairness-adjusted utilitarian, and non-consequentialist. He conceded that two key assumptions must be made. Deploying SRM must not be intrinsically wrong, and "in the presence of climate engineering, the marginal social benefit of further GHG emissions [must be] greater than their marginal social cost". Morrow's analysis concluded that a reduction in mitigation did not necessarily lead to a worse outcome [27]. Bickel and Lane [25] estimated net benefits of various emissions controls strategies with and without the option for climate engineering and found that "it may be economical to replace some degree of emissions reductions with SRM." The study considers both damages from climate change and from mitigation costs.

Empirical evidence for moral hazard in the case of emissions mitigation and climate engineering "can hardly be quantified" [28]. Jesse Reynolds examined several opinion studies and surveys and found that "although each [study] has limitations, all point toward a non-existent or even reverse [climate engineering moral hazard] effect". The results showed that information about government or industry pursuing climate engineering research made people take climate risks more seriously, increasing their desires to mitigate [2].

Though moral hazard in the case of climate engineering is an extremely contentious issue, the evidence does not appear strong enough to eliminate SRM research as an alternative. Robust evidence showing the presence of moral hazard in the case of climate engineering has not been found. Additionally, even if moral hazard effects are present in this case, studies have shown that the outcome would not necessarily be worse. An analysis in chapter five looks at the sensitivity of our decision to the presence of this effect.

4.2 Distributive Justice

Distributive justice objections are concerned with uncertainties about the geographic distribution of harms and benefits from climate engineering. A major claim is that SRM will change regional temperature, precipitation, and climate patterns in a way that threatens the availability of food and water resources to some regions [22]. Simulations of temperature and precipitation patterns have varied results. Some simulations estimate that climate engineering could have positive effects on not only global levels, but also regional levels for both temperature and precipitation patterns, while others conclude the opposite [9, 12, 22]. The uncertainty about the effects of changing regional weather patterns is important to consider in our decision.

Svoboda, Keller, Goes, and Tuana [22] examined distributive justice in the case of climate engineering by critiquing it under the lens of five philosophical theories. Looking at the egalitarian theories of Rawls, Dworkin, Sen, and Arneson as well as desert-based theories of justice, the authors argue that SRM fails to satisfy the requirements of each theory of justice. This study does not provide information about the likelihood of regions or individuals suffering harms, but instead the ethical implications of a possibility of harm. The authors acknowledge that "all realistic responses to climate change…have ethical problems" but do not evaluate other alternatives with the same criteria [22].

A separate study of distributive justice in the case of climate engineering, Wong [30], argues that most distributive justice arguments are focused on outcomes rather than risk. Wong characterizes this as the 'if and then' syndrome and claims it "reminds [of] the danger of overlooking the epistemic gap between 'might' and 'being' (or 'will be') and wrongly emphasizes the merely possible as the most ethically significant". He directly critiques Svoboda [22] for its failure to acknowledge risk, claiming that their arguments are "inconclusive" and "misleading". Rather, the evaluation of distributive justice should focus on sources of risk and preventive measures taken to offset the possible negative consequences in order to better understand a region or individual's likely situation [30].

Distributive justice based objections often combine ethical problems stemming from the technology itself with ethical problems stemming from the way in which the technology is used. What these studies and others fail to acknowledge is that any evaluation of the risk to an individual or region is dependent on the manner and extent to which SRM is implemented and not on the existence of the technology. For example, deploying SRM at different latitudes could limit the effects of particulate matter to certain areas. Additionally if SRM is set to offset different levels of warming, different amounts are necessary and expected damages change accordingly. This idea can be seen in the analysis in Chapter 5.

Another perspective on distributive justice in the case of climate engineering introduces the idea of compensation schemes to the discussion. The Oxford Principles, a set of guidelines to aid in the development of climate engineering governance, question the idea that everyone should benefit equally from research and development of SRM technologies. The principles move beyond the idea of Pareto-optimality even and propose that not only is it acceptable for benefits to vary, but it may also be acceptable that some individuals are made worse off. A study explaining the motivations of the principles refers to the Kaldor-Hicks criterion, a so-called "weaker interpretation" of Pareto optimality in which "some can be rendered worse off provided compensation is in principle payable to them, but [it] does not require that compensation is actually paid" [11].

Though the idea of a compensation scheme appears closely related to the preventive measures mentioned by Wong, he argues that compensation schemes are not very relevant because they are outcome-based in nature. Compensation made after damages from climate engineering has a limited role in the planning process, and compensation made beforehand "does not concern *how* [climate engineering] is to be implemented." In effect, compensation is important but limited in influencing the decision to pursue research [30].

Further difficulties for addressing distributive justice with compensation schemes include the complex international systems needed to organize and administer such a program. An essential concern for any organization implementing a compensation program is liability. Wil Burns proposed that liability may already be included in the preamble to the United Nations Framework Convention on Climate Change (UNFCCC) which states that, "States have … the responsibility to ensure that activities within their jurisdiction or control do not cause damage to the environment of other States…" While certainly plausible, this consideration of liability has not been recognized in the case of climate engineering. Still beyond the problem of questionable liability, he notes difficulties arising from ascribing causality, funding a liability regime, ethical implications of compensating loss of life, and determining offsets for benefits provided by SRM use [14]. Finally, the lack of a precedent for such large transfers of wealth might serve as a final consideration for practicality.

4.3 INTERGENERATIONAL JUSTICE

Intergenerational justice objections are a subcategory of distributive justice claims concerned with uncertainties about the temporal distribution of harms and benefits from climate engineering. Stephen Gardiner argues that the large problem for intergenerational justice is how convenient it is for each current generation to "[exploit] its temporal position," making it possible to "take advantage of the future without the unpleasantness of admitting it" [23]. The previously identified benefit of SRM as a "great experiment" which can be stopped at any time is the primary concern for intergenerational justice objections to climate engineering. If climate engineering were deployed, future generations would experience increased risks of climate damage from rapid warming should the program be terminated [22, 29, 31]. Transferring these risks to future generations fails ethical tests for intergenerational justice [26].

There is uncertainty about the likelihood of a climate engineering program being terminated [22]. Most researchers agree that an unplanned termination of a program could result in more damages to people at the time of termination, but the interpretation of these possible outcomes varies [22, 31]. Svoboda, Keller, Goes, and Tuana [22] examined intergenerational justice in the case of climate engineering in the same manner as done for distributive justice. The study argues that even if future generations are wealthier as a whole, termination of a program would increase inequality and fail to satisfy the requirements for justice provided by the egalitarian theories. Once more the study does not identify the likelihood or extent of future generations suffering but rather treats any possibility of harm as ethically equivalent. Finally, they claim that SRM fails an examination of desert-based justice because future generations who "would not be responsible for having implemented [SRM] do not deserve to be harmed" [22]. This argument is perhaps the strongest ethical objection presented for intergenerational justice. However, the authors fail to acknowledge that the decision to terminate the program (or not to reinstate a terminated program) must be made by future generations. The possible harms which the authors claim as unjust are not features of the technology itself or necessary results of its use; they are the result of a terminated program. That is, the possible harms to future generations are the result of a decision made by future generations, not simply a burden to be carried by uninvolved observers.

Perhaps the authors would still argue that giving future generations the option to make a bad decision is unjust. Many researchers have made arguments for and against this point, commonly referred to as the "arm the future" argument. The arm the future argument is a precautionary measure that claims an "additional option being available is clearly to be preferred to having no choice" [28]. This strategy of adding alternatives works to increase intergenerational justice in two ways. First, it acknowledges that although we have a poor understanding of the preferences of future generations, basic needs can be assumed. In terms of basic needs, having the additional option is better [28]. Second, it acknowledges the risk to future generations following the decision to not pursue climate engineering research. If the additional option to use SRM is not available, future generations face the risk of crossing climate tipping points – risks which could be limited by the SRM option [31].

4.4 Procedural Justice

Procedural justice objections to climate engineering research are largely focused on issues of governance rather than the technology itself. Svoboda, Keller, Goes, and Tuana [22] describe a few theories of procedural justice and claim that "a decision is just if and only if it is reached in the manner it ought to be reached." One major theory discussed claims that a requirement for procedural justice is that everyone has the opportunity to participate in the decision making process. A separate theory of procedural justice requires a public rationale which is related to the decision and subject to appeal. Because SRM is cheap enough to be deployed unilaterally, the authors claim, it fails to meet the requirements of these two theories of justice. It is noted however, that the technology itself may not be procedurally unjust based on this framework for evaluation [22].

4.5 Treating Symptoms

Objections to pursuing climate engineering research frequently include its failure to address the root cause of the problem. It instead treats the symptoms of climate change. This leads to problems in two areas. Climate engineering distracts from treatment of the cause of the original problem, and it does not resolve related problems from the effects of greenhouse gas emissions [7, 12, 22, 26]. The first problem identified is another description of the moral hazard and risk compensation objection and has already been discussed.

The second problem identified, that SRM fails to resolve other problems arising from greenhouse gas emissions, notes that many cycles and weather patterns are not restored by using SRM technologies. For example, ocean acidification continues to impact marine and human ecosystems, and El Niño patterns might change [22]. This argument in effect requires that an alternative address all problems related to climate change. As in the case for moral hazard, society's objectives are key to this objection. If society's objectives are to reduce climate risks and damages and ensure social welfare, there is no need for an alternative to address all carbon emissions related problems. Collier [32] addresses this argument by stating that, "Typically, in an attempt to find a solution to a problem people look to its causes, or yet more fatuously, to its *root* cause. However, there need be no logical connection between the cause of a problem and appropriate or even feasible solutions." With this in mind, other problems caused by greenhouse gas emissions can be addressed by decisions made separately or at a later time.

4.6 Conflict and Termination

The objections to climate engineering research based on conflict and termination bring together many of the previous arguments about justice. Climate change involves the risk of great harms for some countries, but others such as China, Russia and Canada could potentially benefit from a warmer climate. With these differences in benefits, the conflict objection proposes that tensions and conflict will arise in an attempt to control the climate for a country's own benefit [12]. While this may hold true, current military leadership in the United States and the United Kingdom have publicly stated that climate change itself is a threat multiplier for most at risk and instable places in the world. These militaries are focusing on the impacts of the changing climate and how they are impacting geopolitical stability. Projected climate change will add tensions to even developed and stable regions [33]. It is unclear in which way SRM research would differ from this baseline situation in terms of potential for conflict.

Another perspective of the conflict and termination objection considers the possibility of any conflict or the breakdown of an international agreement impeding the continuation of a SRM program, thereby terminating SRM use and leading to abrupt warming. Goes, Tuana, and Keller [26] use this assumption while performing a cost-benefit analysis of climate engineering by including an option for a terminated program. The authors concluded that, with the exception of a small region where SRM causes damages close to 0.5% gross world output and has very low probabilities of intermittency, SRM fails a cost-benefit analysis. The study also notes that such an event also fails tests for distributive, intergenerational, and procedural justice [26]. Bickel and Agrawal [31] responded to Goes, Tuana, and Keller [26] noting that the results of the analysis were "based on their framing of the [SRM]-use decision, rather than on the underlying concept itself." In addition to combining the decision to use SRM with the decision to use a no controls emissions policy, the policy used for comparison when SRM was not used assumed optimal controls. Finally what is most relevant to the conflict and termination objection, the analysis assumed conflict would cause a discontinuation of a SRM program, but did not extend this conflict to the case without SRM use where it would potentially interrupt mitigation efforts. When the analysis is adjusted to place the alternatives in equivalent contexts, Bickel and Agrawal [31] found that "instead of [SRM] failing for any probability of abortment greater than 0.15, the new threshold is 0.89." In consideration of these large differences, it appears the conflict objection is an objection to the governance of the technology rather than to the technology itself.

4.7 Decision to Research and Decision to Use SRM

Until this point in the paper, the decision to pursue research has been considered independent of other decisions. One major objection to pursuing climate engineering research hinges on the idea that this decision is in fact not independent. Once research has been performed, society will definitely use the technology, effectively combining the decision to research and the decision to use SRM. Betz [28] introduces several arguments to this point in an analysis of the arm the future argument. The first argument claims that "the deployment of such technologies will eventually be unavoidable once their development has gained momentum." This in combination with unreflective research lead to development and implementation even when the research shows use is unwarranted. As a final remark it is claimed that once the technology has developed there will be a sort of technological lock-in and that "for this reason, it is easier and safer to stop before getting started than it is to slow development already in progress" [28].

Other large projects with potential risks have also been discussed in a similar manner. Barrett [12] discusses two such projects where apprehension surrounded the development and use of the technologies. A relatable case to climate engineering is that of the Large Hadron Collider at CERN (the European Organization for Nuclear Research). During the planning and construction of this particle collider, there were fears and speculations that such high energy particle collisions could possibly "transform the entire planet Earth into an inert hyperdense sphere about 100m across" or "create a growing black hole…that might destroy not just the Earth but the entire universe" [12]. These claims arose from two large mistakes. The most apparent is the misunderstanding about physical properties of black holes and the potential for a black hole to destroy the universe. The basis for these claims is perhaps more relevant to this discussion; the assumption rested on the idea that such particle collisions had never been observed and the probability of such a reaction could therefore not be estimated. However, collisions at the energy planned for the LHC and even those with higher energy than collisions at the LHC occur naturally in the Earth's atmosphere extremely frequently. If such an event were possible, as proposed by those in objection to the project, it would have been observed already [34]. Somewhat similarly, discussion of the potential risks associated with SRM use often fails to observe that sulfur particles, emitted in fossil fuel power production for example, are already in the atmosphere and presently block nearly 40% of the solar radiation necessary to offset warming [25]. While there are certainly more interactive effects and the issue of scale to consider in the case of SRM, the comparison nonetheless provides insight into prior large-scale projects which were perceived with large potential risks.

The second case mentioned by Barrett [12] discusses the decision to retain or destroy the last known stocks of smallpox. Risk exists in both cases. If the last stocks are destroyed all options for further research or vaccine production are eliminated. However, if the stocks are retained, an accidental or intentional leak could occur, imposing huge losses to the global population which no longer receives vaccinations against the eradicated disease. The last remaining stocks have been held at World Health Organization centers in Atlanta and Moscow since the eradication of smallpox in 1977. The stocks have not been used inappropriately and both the United States and Russia have argued to retain the stocks to manage risks in the unprecedented event of smallpox-based biological warfare [12]. Looking at this case under the lens of the objections presented in Betz [28], neither development of the stock nor research after the fact led to deployment or inappropriate use of the disease stock. It appears then that the combination of the decision to develop and the decision to use a technology might actually be an uncertainty rather than a certain coupling. This proposal is further considered in the case of the decision to pursue climate engineering research in the next chapter.

Chapter 5: Decision Sensitivity

This chapter presents an original analysis of the sensitivity of our decision to pursue climate engineering research to the probability of key ethical objections. Bickel [18] developed a model which adds a climate engineering option to the Nordhaus [8] DICE model (the Dynamically Integrated Climate-Economy model). The model considers economic and geophysical relationships over the course of 200 years and calculates the present value of different policy alternatives such as no emissions controls, optimal emissions controls, and temperature limits [8]. The analysis in this report is based on outputs from the Bickel [18] model.

In this analysis the assumptions of two main ethical objections to climate engineering research, that climate engineering research would lessen society's mitigation efforts and would necessarily lead to deployment of the technology, are incorporated into the outputs of the Bickel [18] model to see how the decision to research changes with the probability of each assumption. The decision is analyzed across two different levels of SRM use, first at a level which uses variable amounts of SRM in each period to completely offset any increase in the global average temperature, and then at a level which also uses variable amounts of SRM but allows up to a 2°C increase in the global average temperature.

For each level of SRM use there are two sets of emissions controls policies compared: in the first there is a trade-off between optimal emissions controls and no controls and in the second there is a trade-off between an emissions control policy to limit temperature increase to 2°C (L2C) and no controls. The no controls policy is also referred to as a business-as-usual policy and acts as a baseline. It does not imply that society will stop all current efforts to limit emissions, but rather it follows the international trend up through 2007. The optimal controls policy solves for the economically efficient policy, balancing current abatement costs and future damages from warming. The L2C policy sets emissions at a level expected to limit temperature increases to a 2°C increase from 1900 levels [8]. These comparison sets are intentionally set in pairs of two to emphasize the effects of an individual policy rather than the effects of more complex changes to a portfolio with multiple options.

Finally each combination of emissions controls policies and SRM levels are compared across six levels of damages caused by SRM, with each level measured as a percent of global world output. This assumes that after SRM research is conducted researchers will know how much damage is caused by the technology. Rather than a decision to deploy the technology at this point, the analysis considers the deployment of the technology an uncertainty to evaluate the significance of this objection to research to society's decision.

The result is a table which shows at which probabilities of deploying the damaging technology the decision to research would still provide net benefits to society. Along the other axis of each table, the probability of achieving the desired emissions controls goal is adjusted to show how its changes affect the decision to research. This portion of the analysis looks at how the trade-off between emissions controls policies hypothesized in the moral hazard and risk compensation objections to research would change the net benefits provided by researching SRM. The remainder of this chapter looks at each scenario described above. The sub-sections begin with a decision tree to illustrate the scenario and are followed by a description of the scenario, the data table outputs for each level of SRM damage, and a description of what is implied by the output tables.

5.1 SRM Level Complete Offset

5.1.1 Optimal Emissions Controls

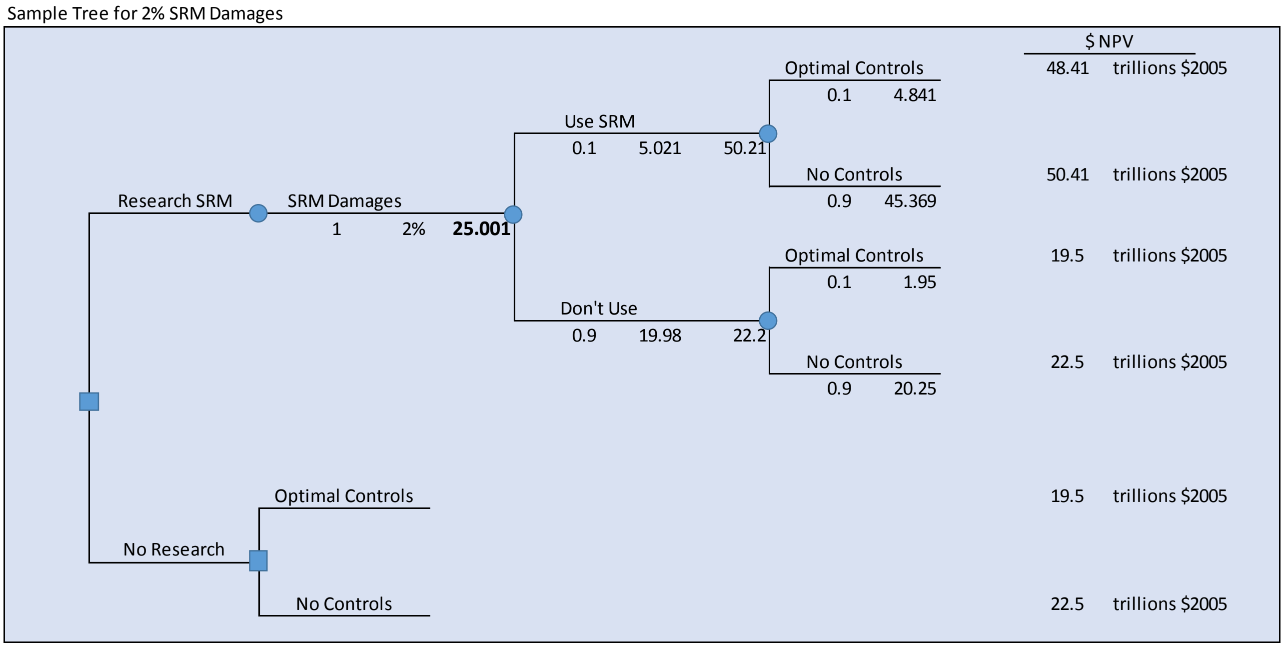

Figure 1: Sample Tree for 2% SRM Damages when SRM Level is Complete Offset with Optimal Emissions Controls

In the following two scenarios SRM is used until it completely offsets any temperature changes. The decision tree above shows the sample case when research for SRM reveals that there will be certain damages of 2% of gross world output. The model assumes there is uncertainty as to whether SRM will be deployed or not. This is indicated in the braches marked “Use SRM” and “Don’t Use”. In each case there is also uncertainty about the emissions control method society will adopt after pursuing research. These are noted by the “Optimal Controls” and “No Controls” branches on the right of the tree. The probability of each control method is independent of whether SRM is deployed or not, because the underlying assumption in the moral hazard and risk compensation argument is that research alone detracts from the emissions control method.

Finally, the right-most column shows the NPV for each scenario. These values are outputs of the Bickel-DICE modification, which includes an option for climate engineering. For this exercise, the emissions control method, SRM level, and SRM damages were adjusted as inputs. The resulting output is shown as the NPV of damages in trillions of 2005 dollars.

The following pages show the decision to pursue research when SRM will be used at levels to completely offset warming. Each page is a comparison to an individual emissions control strategy that is a sub-decision of the "No Research" branch. SRM damages ranging from 0% to 5% of gross world output are considered in each comparison. As a benchmark for the range of these damages, Nordhaus [8] estimates that "the economic damages from climate change with no interventions will be on the order of 2.5 percent of world output per year." Moving from left to right across each table, the probability of optimal controls increases by 0.1 per cell. Moving from the top to the bottom of the table, the probability of deploying SRM increases by 0.1 per cell. In the green cells, the NPV of damages for pursuing research is less than or equal to the NPV of damages for the emissions control strategy in comparison. Because this is a comparison of damages, a lower NPV indicates a good decision. In the red cells, the NPV of damages for pursuing research is higher than the NPV of damages for the emissions control strategy in comparison, indicating a bad decision.

Emission Control Level No Controls

This page shows our decision when SRM is used to completely offset warming to compared to a no controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research includes optimal controls and no controls.

Figure 2: Model Output Tables for 0% - 5% Damages from SRM with No Controls Baseline

The first table shows the results when SRM has damages amounting to 0% of gross world output. In this case, researching SRM provides greater benefits than the no research alternative for all probabilities of achieving optimal controls, noted by ‘Q (OC)’ in the table, and for all probabilities of deploying the technology after research, noted by ‘P (Use SRM)’ in the table. These benefits fade very quickly however, as illustrated by the differences in output already in the tables for 1% and 2% damages caused by SRM. In the case where SRM causes 1% damages, there is a sizable region where the decision to pursue research still has greater benefits than no research. This area describes situations where the probability of optimal emissions is high and the probability of deploying the technology is low. In later cases with 2% damages and even more obviously with 3% - 5% damages, researching SRM only provides equivalent or larger benefits than no research when the technology will not be used. Because the baseline emissions policy considered when research is not conducted is no controls, the probability of moral hazard does not change these results as damages continue to increase from 3% damages onward.

Emission Control Level Optimal Controls

This page shows our decision when SRM is used to completely offset warming compared to an optimal controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research again includes optimal controls and no controls.

0% SRM Damages vs No SRM+Optimal Controls

P (Use SRM) \ Q (OC)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-1.0	-0.8	-0.5	-0.3	0.0	0.2	0.5	0.7	1.0	1.2	1.5
0.2	0.9	1.1	1.4	1.6	1.8	2.0	2.2	2.4	2.6	2.8	3.0
0.3	2.9	3.1	3.2	3.4	3.5	3.7	3.8	4.0	4.1	4.3	4.4
0.4	4.9	5.0	5.1	5.2	5.3	5.4	5.5	5.6	5.7	5.8	5.9
0.5	6.9	6.9	7.0	7.0	7.1	7.1	7.2	7.3	7.3	7.4	7.4
0.6	8.8	8.8	8.9	8.9	8.9	8.9	8.9	8.9	8.9	8.9	8.9
0.7	10.8	10.8	10.7	10.7	10.6	10.6	10.6	10.5	10.5	10.4	10.4
0.8	12.8	12.7	12.6	12.5	12.4	12.3	12.2	12.1	12.0	12.0	11.9
0.9	14.8	14.6	14.5	14.3	14.2	14.1	13.9	13.8	13.6	13.5	13.3
1.0	16.7	16.5	16.4	16.2	16.0	15.8	15.6	15.4	15.2	15.0	14.8

1% SRM Damages vs No SRM+Optimal Controls

P (Use SRM) \ Q (OC)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-3.4	-3.1	-2.9	-2.6	-2.3	-2.1	-1.8	-1.5	-1.2	-1.0	-0.7
0.2	-3.8	-3.6	-3.3	-3.1	-2.9	-2.6	-2.4	-2.1	-1.9	-1.7	-1.4
0.3	-4.2	-4.0	-3.8	-3.6	-3.4	-3.2	-3.0	-2.8	-2.5	-2.3	-2.1
0.4	-4.6	-4.5	-4.3	-4.1	-3.9	-3.7	-3.6	-3.4	-3.2	-3.0	-2.8
0.5	-5.0	-4.9	-4.7	-4.6	-4.4	-4.3	-4.1	-4.0	-3.8	-3.7	-3.5
0.6	-5.5	-5.3	-5.2	-5.1	-5.0	-4.9	-4.7	-4.6	-4.5	-4.4	-4.3
0.7	-5.9	-5.8	-5.7	-5.6	-5.5	-5.4	-5.3	-5.2	-5.1	-5.1	-5.0
0.8	-6.3	-6.2	-6.2	-6.1	-6.0	-6.0	-5.9	-5.9	-5.8	-5.7	-5.7
0.9	-6.7	-6.7	-6.6	-6.6	-6.6	-6.5	-6.5	-6.5	-6.4	-6.4	-6.4
1.0	-7.1	-7.1	-7.1	-7.1	-7.1	-7.1	-7.1	-7.1	-7.1	-7.1	-7.1

2% SRM Damages vs No SRM+Optimal Controls

P (Use SRM) \ Q (OC)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-5.8	-5.5	-5.2	-4.9	-4.6	-4.3	-4.1	-3.8	-3.5	-3.2	-2.9
0.2	-8.6	-8.3	-8.0	-7.7	-7.5	-7.2	-6.9	-6.6	-6.3	-6.1	-5.8
0.3	-11.4	-11.1	-10.8	-10.6	-10.3	-10.0	-9.8	-9.5	-9.2	-8.9	-8.7
0.4	-14.2	-13.9	-13.6	-13.4	-13.1	-12.9	-12.6	-12.3	-12.1	-11.8	-11.6
0.5	-17.0	-16.7	-16.5	-16.2	-16.0	-15.7	-15.5	-15.2	-15.0	-14.7	-14.5
0.6	-19.7	-19.5	-19.3	-19.0	-18.8	-18.5	-18.3	-18.1	-17.8	-17.6	-17.3
0.7	-22.5	-22.3	-22.1	-21.8	-21.6	-21.4	-21.2	-20.9	-20.7	-20.5	-20.2
0.8	-25.3	-25.1	-24.9	-24.7	-24.4	-24.2	-24.0	-23.8	-23.6	-23.3	-23.1
0.9	-28.1	-27.9	-27.7	-27.5	-27.3	-27.1	-26.9	-26.6	-26.4	-26.2	-26.0
1.0	-30.9	-30.7	-30.5	-30.3	-30.1	-29.9	-29.7	-29.5	-29.3	-29.1	-28.9

3% SRM Damages vs No SRM+Optimal Controls

P (Use SRM) \ Q (OC)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-8.2	-7.9	-7.6	-7.3	-7.0	-6.6	-6.3	-6.0	-5.7	-5.4	-5.1
0.2	-13.4	-13.1	-12.7	-12.4	-12.1	-11.8	-11.5	-11.1	-10.8	-10.5	-10.2
0.3	-18.6	-18.2	-17.9	-17.6	-17.3	-16.9	-16.6	-16.3	-15.9	-15.6	-15.3
0.4	-23.8	-23.4	-23.1	-22.7	-22.4	-22.1	-21.7	-21.4	-21.0	-20.7	-20.3
0.5	-29.0	-28.6	-28.3	-27.9	-27.6	-27.2	-26.8	-26.5	-26.1	-25.8	-25.4
0.6	-34.2	-33.8	-33.4	-33.1	-32.7	-32.3	-32.0	-31.6	-31.2	-30.9	-30.5
0.7	-39.4	-39.0	-38.6	-38.2	-37.9	-37.5	-37.1	-36.7	-36.3	-36.0	-35.6
0.8	-44.6	-44.2	-43.8	-43.4	-43.0	-42.6	-42.2	-41.8	-41.4	-41.1	-40.7
0.9	-49.7	-49.3	-48.9	-48.5	-48.2	-47.8	-47.4	-47.0	-46.6	-46.2	-45.8
1.0	-64.9	-64.5	-64.1	-63.7	-63.3	-62.9	-62.5	-62.1	-61.7	-61.3	-60.8

4% SRM Damages vs No SRM+Optimal Controls

P (Use SRM) \ Q (OC)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-10.6	-10.3	-9.9	-9.6	-9.3	-8.9	-8.6	-8.3	-8.0	-7.6	-7.3
0.2	-18.2	-17.8	-17.5	-17.1	-16.8	-16.4	-16.0	-15.7	-15.3	-14.9	-14.6
0.3	-25.8	-25.4	-25.0	-24.6	-24.2	-23.8	-23.4	-23.0	-22.7	-22.3	-21.9
0.4	-33.4	-33.0	-32.6	-32.1	-31.7	-31.3	-30.9	-30.4	-30.0	-29.6	-29.1
0.5	-41.0	-40.6	-40.1	-39.7	-39.2	-38.7	-38.3	-37.8	-37.4	-36.9	-36.4
0.6	-48.6	-48.1	-47.7	-47.2	-46.7	-46.2	-45.7	-45.2	-44.7	-44.2	-43.7
0.7	-56.2	-55.7	-55.2	-54.7	-54.1	-53.6	-53.1	-52.6	-52.1	-51.5	-51.0
0.8	-63.8	-63.3	-62.7	-62.2	-61.6	-61.1	-60.5	-60.0	-59.4	-58.8	-58.3
0.9	-71.5	-70.9	-70.3	-69.7	-69.1	-68.5	-67.9	-67.3	-66.8	-66.2	-65.6
1.0	-79.1	-78.4	-77.8	-77.2	-76.6	-76.0	-75.3	-74.7	-74.1	-73.5	-72.9

5% SRM Damages vs No SRM+Optimal Controls

P (Use SRM) \ Q (OC)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-13.0	-12.7	-12.3	-12.0	-11.6	-11.3	-10.9	-10.6	-10.2	-9.8	-9.5
0.2	-23.1	-22.6	-22.2	-21.8	-21.4	-21.0	-20.6	-20.2	-19.8	-19.4	-19.0
0.3	-33.1	-32.6	-32.2	-31.7	-31.2	-30.8	-30.3	-29.9	-29.4	-28.9	-28.5
0.4	-43.1	-42.6	-42.1	-41.6	-41.1	-40.5	-40.0	-39.5	-39.0	-38.5	-38.0
0.5	-53.1	-52.6	-52.0	-51.4	-50.9	-50.3	-49.7	-49.2	-48.6	-48.0	-47.4
0.6	-63.2	-62.5	-61.9	-61.3	-60.7	-60.1	-59.4	-58.8	-58.2	-57.6	-56.9
0.7	-73.2	-72.5	-71.8	-71.2	-70.5	-69.8	-69.1	-68.5	-67.8	-67.1	-66.4
0.8	-83.2	-82.5	-81.8	-81.0	-80.3	-79.6	-78.8	-78.1	-77.4	-76.6	-75.9
0.9	-93.3	-92.5	-91.7	-90.9	-90.1	-89.3	-88.5	-87.8	-87.0	-86.2	-85.4
1.0	-103.3	-102.4	-101.6	-100.8	-99.9	-99.1	-98.2	-97.4	-96.6	-95.7	-94.9

Figure 3: Model Output Tables for 0% - 5% Damages from SRM with Optimal Controls Baseline

The first table again shows the results when SRM damages are 0% of gross output. In this case research provides more benefits than the no research alternative for most probabilities of deployment after research and achieving optimal emissions controls. The small area where this differs occurs because the probability of SRM use is too low to compensate for the damages incurred by the decreased probability of achieving optimal emissions. The higher damages of research are quickly overcome with as little as a 20% probability of SRM use. However, as soon as damages from SRM are introduced, no research provides greater benefits than research for each probability of using SRM and each probability of achieving optimal emissions except the case where the probability of use is zero and the probability of optimal controls is one.

Considering that the baseline for comparison when there is no research in this case is optimal controls, these results are somewhat intuitive. The optimal controls are by definition the most economically efficient means of mitigation, so an alternative with any detraction from certainly achieving optimal controls would make the alternative worse. This leaves only the rightmost column in the output table, where the probability of achieving optimal controls is one, as possible cases for research having greater benefits than no research. However, once damages have been introduced for SRM use, as they have been in the tables for 1% - 5% damages from SRM use, any probability of using the technology also lowers the benefits when compared to optimal controls. The only remaining case where research has equivalent benefits to the no research alternative is then the case described previously where the probability of use is zero and the probability of optimal controls is one.

5.1.2 Limit 2°C Emissions Controls

Sample Tree for 2% SRM Damages

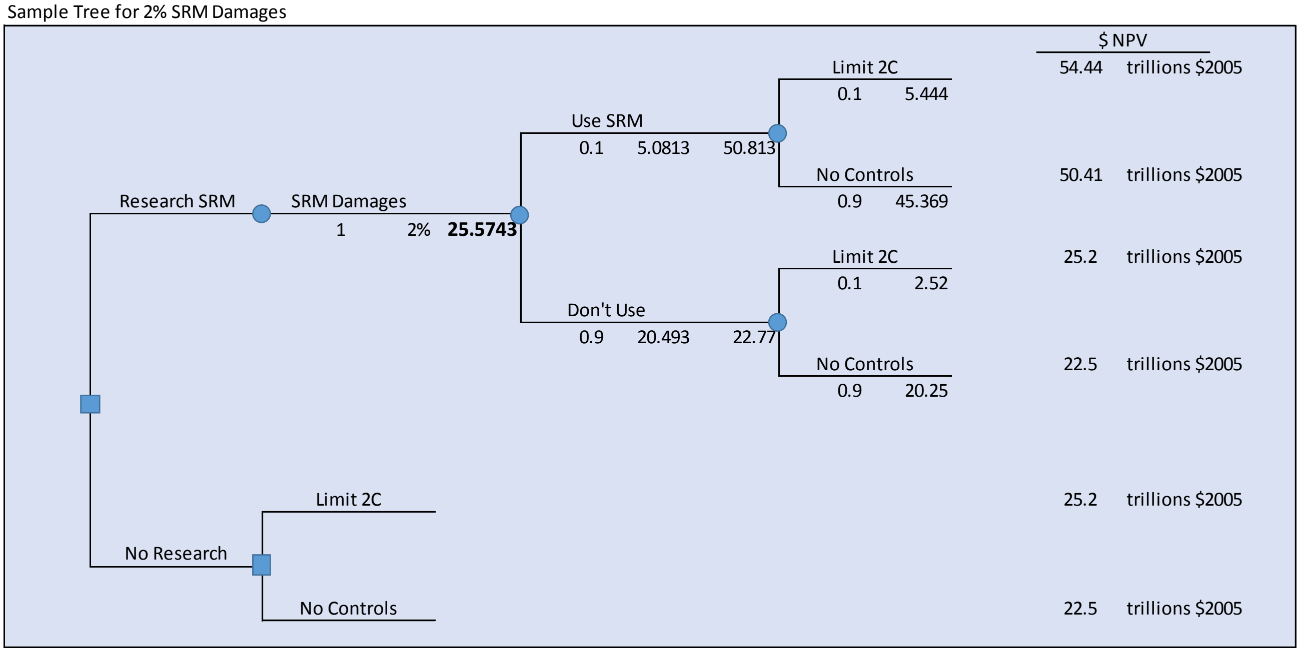

Figure 4: Sample Tree for 2% SRM Damages when SRM Level is Complete Offset with Limit 2°C Emissions Controls

In the following two scenarios SRM is again used until it completely offsets any temperature changes. The decision tree above shows the sample case when research for SRM reveals that there will be certain damages of 2% of gross world output. The uncertainties and assumptions for the model are the same as in the previous scenarios. Rather than looking at an optimal emissions control method though, these scenarios compare the no controls method to an emissions control method which limits temperature change to 2°C as noted by the "Limit 2C" branch in the decision tree. Once more, moving from left to right across each table, the probability of optimal controls increases by 0.1 per cell. Moving from the top to the bottom of the table, the probability of deploying SRM increases by 0.1 per cell. In the green cells, the NPV of damages for pursuing research is less than or equal to the NPV of damages for the emissions control strategy in comparison.

Emission Control Level No Controls

This page shows our decision when SRM is used to completely offset warming compared to a no controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research includes L2C controls and no controls.

0% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	2.0	1.6	1.3	0.9	0.5	0.2	-0.2	-0.5	-0.9	-1.2	-1.6
0.2	3.9	3.5	3.1	2.6	2.2	1.7	1.3	0.8	0.4	0.0	-0.5
0.3	5.9	5.4	4.9	4.3	3.8	3.3	2.7	2.2	1.7	1.2	0.6
0.4	7.9	7.3	6.7	6.0	5.4	4.8	4.2	3.6	3.0	2.4	1.7
0.5	9.9	9.2	8.5	7.8	7.1	6.4	5.7	5.0	4.3	3.5	2.8
0.6	11.8	11.1	10.3	9.5	8.7	7.9	7.1	6.3	5.5	4.7	4.0
0.7	13.8	12.9	12.1	11.2	10.3	9.4	8.6	7.7	6.8	5.9	5.1
0.8	15.8	14.8	13.9	12.9	11.9	11.0	10.0	9.1	8.1	7.1	6.2
0.9	17.8	16.7	15.7	14.6	13.6	12.5	11.5	10.4	9.4	8.3	7.3
1.0	19.7	18.6	17.5	16.3	15.2	14.1	12.9	11.8	10.7	9.5	8.4

1% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-0.4	-0.7	-1.0	-1.4	-1.7	-2.0	-2.3	-2.6	-3.0	-3.3	-3.6
0.2	-0.8	-1.2	-1.6	-1.9	-2.3	-2.7	-3.0	-3.4	-3.8	-4.1	-4.5
0.3	-1.2	-1.6	-2.1	-2.5	-2.9	-3.3	-3.7	-4.2	-4.6	-5.0	-5.4
0.4	-1.6	-2.1	-2.6	-3.0	-3.5	-4.0	-4.4	-4.9	-5.4	-5.8	-6.3
0.5	-2.0	-2.6	-3.1	-3.6	-4.1	-4.6	-5.1	-5.7	-6.2	-6.7	-7.2
0.6	-2.5	-3.0	-3.6	-4.2	-4.7	-5.3	-5.9	-6.4	-7.0	-7.6	-8.1
0.7	-2.9	-3.5	-4.1	-4.7	-5.3	-5.9	-6.6	-7.2	-7.8	-8.4	-9.0
0.8	-3.3	-3.9	-4.6	-5.3	-5.9	-6.6	-7.3	-7.9	-8.6	-9.3	-9.9
0.9	-3.7	-4.4	-5.1	-5.8	-6.5	-7.3	-8.0	-8.7	-9.4	-10.1	-10.8
1.0	-4.1	-4.9	-5.6	-6.4	-7.1	-7.9	-8.7	-9.4	-10.2	-11.0	-11.7

2% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-2.8	-3.1	-3.4	-3.6	-3.9	-4.2	-4.5	-4.8	-5.1	-5.3	-5.6
0.2	-5.6	-5.9	-6.2	-6.5	-6.8	-7.1	-7.4	-7.7	-8.0	-8.3	-8.5
0.3	-8.4	-8.7	-9.0	-9.3	-9.6	-9.9	-10.2	-10.5	-10.9	-11.2	-11.5
0.4	-11.2	-11.5	-11.8	-12.1	-12.5	-12.8	-13.1	-13.4	-13.7	-14.1	-14.4
0.5	-14.0	-14.3	-14.6	-15.0	-15.3	-15.6	-16.0	-16.3	-16.6	-17.0	-17.3
0.6	-16.7	-17.1	-17.4	-17.8	-18.1	-18.5	-18.8	-19.2	-19.5	-19.9	-20.2
0.7	-19.5	-19.9	-20.3	-20.6	-21.0	-21.4	-21.7	-22.1	-22.4	-22.8	-23.2
0.8	-22.3	-22.7	-23.1	-23.5	-23.8	-24.2	-24.6	-25.0	-25.3	-25.7	-26.1
0.9	-25.1	-25.5	-25.9	-26.3	-26.7	-27.1	-27.5	-27.8	-28.2	-28.6	-29.0
1.0	-27.9	-28.3	-28.7	-29.1	-29.5	-29.9	-30.3	-30.7	-31.1	-31.5	-31.9

3% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-5.2	-5.4	-5.7	-5.9	-6.2	-6.4	-6.7	-6.9	-7.1	-7.4	-7.6
0.2	-10.4	-10.6	-10.8	-11.0	-11.3	-11.5	-11.7	-11.9	-12.1	-12.4	-12.6
0.3	-15.6	-15.8	-16.0	-16.2	-16.4	-16.5	-16.7	-16.9	-17.1	-17.3	-17.5
0.4	-20.8	-20.9	-21.1	-21.3	-21.4	-21.6	-21.8	-21.9	-22.1	-22.3	-22.4
0.5	-26.0	-26.1	-26.3	-26.4	-26.5	-26.7	-26.8	-27.0	-27.1	-27.2	-27.4
0.6	-31.2	-31.3	-31.4	-31.5	-31.6	-31.7	-31.9	-32.0	-32.1	-32.2	-32.3
0.7	-36.4	-36.4	-36.5	-36.6	-36.7	-36.8	-36.9	-37.0	-37.1	-37.2	-37.3
0.8	-41.6	-41.6	-41.7	-41.7	-41.8	-41.9	-41.9	-42.0	-42.1	-42.1	-42.2
0.9	-46.7	-46.8	-46.8	-46.9	-46.9	-46.9	-47.0	-47.0	-47.0	-47.1	-47.1
1.0	-51.9	-52.0	-52.0	-52.0	-52.0	-52.0	-52.0	-52.0	-52.0	-52.0	-52.1

4% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-7.6	-7.8	-8.0	-8.2	-8.4	-8.6	-8.8	-9.0	-9.2	-9.4	-9.6
0.2	-15.2	-15.4	-15.5	-15.6	-15.8	-15.9	-16.0	-16.2	-16.3	-16.5	-16.6
0.3	-22.8	-22.9	-23.0	-23.0	-23.1	-23.2	-23.3	-23.3	-23.4	-23.5	-23.5
0.4	-30.4	-30.4	-30.4	-30.4	-30.5	-30.5	-30.5	-30.5	-30.5	-30.5	-30.5
0.5	-38.0	-38.0	-37.9	-37.9	-37.8	-37.7	-37.7	-37.6	-37.6	-37.5	-37.4
0.6	-45.6	-45.5	-45.4	-45.3	-45.1	-45.0	-44.9	-44.8	-44.6	-44.5	-44.4
0.7	-53.2	-53.1	-52.9	-52.7	-52.5	-52.3	-52.1	-51.9	-51.7	-51.5	-51.3
0.8	-60.8	-60.6	-60.3	-60.1	-59.8	-59.6	-59.3	-59.1	-58.8	-58.5	-58.3
0.9	-68.5	-68.1	-67.8	-67.5	-67.2	-66.8	-66.5	-66.2	-65.9	-65.6	-65.2
1.0	-76.1	-75.7	-75.3	-74.9	-74.5	-74.1	-73.7	-73.3	-73.0	-72.6	-72.2

5% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-10.0	-10.2	-10.4	-10.5	-10.7	-10.8	-11.0	-11.2	-11.3	-11.5	-11.7
0.2	-20.1	-20.1	-20.2	-20.2	-20.3	-20.3	-20.4	-20.5	-20.5	-20.6	-20.6
0.3	-30.1	-30.0	-30.0	-29.9	-29.9	-29.8	-29.8	-29.7	-29.7	-29.6	-29.6
0.4	-40.1	-40.0	-39.8	-39.6	-39.5	-39.3	-39.2	-39.0	-38.9	-38.7	-38.5
0.5	-50.1	-49.9	-49.6	-49.4	-49.1	-48.8	-48.6	-48.3	-48.0	-47.8	-47.5
0.6	-60.2	-59.8	-59.4	-59.1	-58.7	-58.3	-57.9	-57.6	-57.2	-56.8	-56.5
0.7	-70.2	-69.7	-69.2	-68.8	-68.3	-67.8	-67.3	-66.9	-66.4	-65.9	-65.4
0.8	-80.2	-79.6	-79.1	-78.5	-77.9	-77.3	-76.7	-76.1	-75.6	-75.0	-74.4
0.9	-90.3	-89.6	-88.9	-88.2	-87.5	-86.8	-86.1	-85.4	-84.7	-84.0	-83.3
1.0	-100.3	-99.5	-98.7	-97.9	-97.1	-96.3	-95.5	-94.7	-93.9	-93.1	-92.3

Figure 5: Model Output Tables for 0% - 5% Damages from SRM with No Controls Baseline

The first table shows the results when SRM has damages amounting to 0% of gross world output. In this case, researching SRM provides greater benefits than no research for most probabilities of achieving optimal controls and for most probabilities of deploying the technology after research. There is a small area in this first table where research provides fewer benefits than the no research alternative. This occurs when there are low probabilities of using the technology and high probabilities of achieving L2C controls. The damages from a L2C emissions control method are actually greater than those of a no controls policy because so much money is spent on mitigation efforts. Additionally, because the level of SRM use is set to completely offset warming and the L2C controls limit temperature change to 2°C by definition, damages from temperature increases are smaller when the technology is used. For this reason, the region with high probability of achieving L2C controls and low probability of SRM use results in greater damages.

As in the previous case where optimal controls and no controls were trade-offs, the benefits of research fade very quickly. Once damages are introduced for using SRM, as shown in the tables for 1% - 5% damages caused by SRM, any probability of using the technology also lowers the benefits when compared to no controls. The only remaining case where research does not provide more damages than no research is the case where the probability of use is zero and the probability of L2C controls is also zero.

Emission Control Level Limit 2°C

This page shows our decision when SRM is used to completely offset warming compared to an emissions scenario which limits temperature increase to 2°C. The trade-off set for emissions controls after research includes L2C controls and no controls.

0% SRM Damages vs No SRM+ L2C Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	4.7	4.3	4.0	3.6	3.2	2.9	2.5	2.2	1.8	1.5	1.1
0.2	6.6	6.2	5.8	5.3	4.9	4.4	4.0	3.5	3.1	2.7	2.2
0.3	8.6	8.1	7.6	7.0	6.5	6.0	5.4	4.9	4.4	3.9	3.3
0.4	10.6	10.0	9.4	8.7	8.1	7.5	6.9	6.3	5.7	5.1	4.4
0.5	12.6	11.9	11.2	10.5	9.8	9.1	8.4	7.7	7.0	6.2	5.5
0.6	14.5	13.8	13.0	12.2	11.4	10.6	9.8	9.0	8.2	7.4	6.7
0.7	16.5	15.6	14.8	13.9	13.0	12.1	11.3	10.4	9.5	8.6	7.8
0.8	18.5	17.5	16.6	15.6	14.6	13.7	12.7	11.8	10.8	9.8	8.9
0.9	20.5	19.4	18.4	17.3	16.3	15.2	14.2	13.1	12.1	11.0	10.0
1.0	22.4	21.3	20.2	19.0	17.9	16.8	15.6	14.5	13.4	12.2	11.1

1% SRM Damages vs No SRM +L2C Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	2.3	2.0	1.7	1.3	1.0	0.7	0.4	0.1	-0.3	-0.6	-0.9
0.2	1.9	1.5	1.1	0.8	0.4	0.0	-0.3	-0.7	-1.1	-1.4	-1.8
0.3	1.5	1.1	0.6	0.2	-0.2	-0.6	-1.0	-1.5	-1.9	-2.3	-2.7
0.4	1.1	0.6	0.1	-0.3	-0.8	-1.3	-1.7	-2.2	-2.7	-3.1	-3.6
0.5	0.7	0.1	-0.4	-0.9	-1.4	-1.9	-2.4	-3.0	-3.5	-4.0	-4.5
0.6	0.2	-0.3	-0.9	-1.5	-2.0	-2.6	-3.2	-3.7	-4.3	-4.9	-5.4
0.7	-0.2	-0.8	-1.4	-2.0	-2.6	-3.2	-3.9	-4.5	-5.1	-5.7	-6.3
0.8	-0.6	-1.2	-1.9	-2.6	-3.2	-3.9	-4.6	-5.2	-5.9	-6.6	-7.2
0.9	-1.0	-1.7	-2.4	-3.1	-3.8	-4.6	-5.3	-6.0	-6.7	-7.4	-8.1
1.0	-1.4	-2.2	-2.9	-3.7	-4.4	-5.2	-6.0	-6.7	-7.5	-8.3	-9.0

2% SRM Damages vs No SRM+ L2C Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	-0.1	-0.4	-0.7	-0.9	-1.2	-1.5	-1.8	-2.1	-2.4	-2.6	-2.9
0.2	-2.9	-3.2	-3.5	-3.8	-4.1	-4.4	-4.7	-5.0	-5.3	-5.6	-5.8
0.3	-5.7	-6.0	-6.3	-6.6	-6.9	-7.2	-7.5	-7.8	-8.2	-8.5	-8.8
0.4	-8.5	-8.8	-9.1	-9.4	-9.8	-10.1	-10.4	-10.7	-11.0	-11.4	-11.7
0.5	-11.3	-11.6	-11.9	-12.3	-12.6	-12.9	-13.3	-13.6	-13.9	-14.3	-14.6
0.6	-14.0	-14.4	-14.7	-15.1	-15.4	-15.8	-16.1	-16.5	-16.8	-17.2	-17.5
0.7	-16.8	-17.2	-17.6	-17.9	-18.3	-18.7	-19.0	-19.4	-19.7	-20.1	-20.5
0.8	-19.6	-20.0	-20.4	-20.8	-21.1	-21.5	-21.9	-22.3	-22.6	-23.0	-23.4
0.9	-22.4	-22.8	-23.2	-23.6	-24.0	-24.4	-24.8	-25.1	-25.5	-25.9	-26.3
1.0	-25.2	-25.6	-26.0	-26.4	-26.8	-27.2	-27.6	-28.0	-28.4	-28.8	-29.2

3% SRM Damages vs No SRM+ L2C Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	-2.5	-2.7	-3.0	-3.2	-3.5	-3.7	-4.0	-4.2	-4.4	-4.7	-4.9
0.2	-7.7	-7.9	-8.1	-8.3	-8.6	-8.8	-9.0	-9.2	-9.4	-9.7	-9.9
0.3	-12.9	-13.1	-13.3	-13.5	-13.7	-13.8	-14.0	-14.2	-14.4	-14.6	-14.8
0.4	-18.1	-18.2	-18.4	-18.6	-18.7	-18.9	-19.1	-19.2	-19.4	-19.6	-19.7
0.5	-23.3	-23.4	-23.6	-23.7	-23.8	-24.0	-24.1	-24.3	-24.4	-24.5	-24.7
0.6	-28.5	-28.6	-28.7	-28.8	-28.9	-29.0	-29.2	-29.3	-29.4	-29.5	-29.6
0.7	-33.7	-33.7	-33.8	-33.9	-34.0	-34.1	-34.2	-34.3	-34.4	-34.5	-34.6
0.8	-38.9	-38.9	-39.0	-39.0	-39.1	-39.2	-39.2	-39.3	-39.4	-39.4	-39.5
0.9	-44.0	-44.1	-44.1	-44.2	-44.2	-44.2	-44.3	-44.3	-44.3	-44.4	-44.4
1.0	-49.2	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.4

3% SRM Damages vs No SRM+ L2C Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	-2.5	-2.7	-3.0	-3.2	-3.5	-3.7	-4.0	-4.2	-4.4	-4.7	-4.9
0.2	-7.7	-7.9	-8.1	-8.3	-8.6	-8.8	-9.0	-9.2	-9.4	-9.7	-9.9
0.3	-12.9	-13.1	-13.3	-13.5	-13.7	-13.8	-14.0	-14.2	-14.4	-14.6	-14.8
0.4	-18.1	-18.2	-18.4	-18.6	-18.7	-18.9	-19.1	-19.2	-19.4	-19.6	-19.7
0.5	-23.3	-23.4	-23.6	-23.7	-23.8	-24.0	-24.1	-24.3	-24.4	-24.5	-24.7
0.6	-28.5	-28.6	-28.7	-28.8	-28.9	-29.0	-29.2	-29.3	-29.4	-29.5	-29.6
0.7	-33.7	-33.7	-33.8	-33.9	-34.0	-34.1	-34.2	-34.3	-34.4	-34.5	-34.6
0.8	-38.9	-38.9	-39.0	-39.0	-39.1	-39.2	-39.2	-39.3	-39.4	-39.4	-39.5
0.9	-44.0	-44.1	-44.1	-44.2	-44.2	-44.2	-44.3	-44.3	-44.3	-44.4	-44.4
1.0	-49.2	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.3	-49.4

5% SRM Damages vs No SRM+ L2C Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	-7.3	-7.5	-7.7	-7.8	-8.0	-8.1	-8.3	-8.5	-8.6	-8.8	-9.0
0.2	-17.4	-17.4	-17.5	-17.5	-17.6	-17.6	-17.7	-17.8	-17.8	-17.9	-17.9
0.3	-27.4	-27.3	-27.3	-27.2	-27.2	-27.1	-27.1	-27.0	-27.0	-26.9	-26.9
0.4	-37.4	-37.3	-37.1	-36.9	-36.8	-36.6	-36.5	-36.3	-36.2	-36.0	-35.8
0.5	-47.4	-47.2	-46.9	-46.7	-46.4	-46.1	-45.9	-45.6	-45.3	-45.1	-44.8
0.6	-57.5	-57.1	-56.7	-56.4	-56.0	-55.6	-55.2	-54.9	-54.5	-54.1	-53.8
0.7	-67.5	-67.0	-66.5	-66.1	-65.6	-65.1	-64.6	-64.2	-63.7	-63.2	-62.7
0.8	-77.5	-76.9	-76.4	-75.8	-75.2	-74.6	-74.0	-73.4	-72.9	-72.3	-71.7
0.9	-87.6	-86.9	-86.2	-85.5	-84.8	-84.1	-83.4	-82.7	-82.0	-81.3	-80.6
1.0	-97.6	-96.9	-96.0	-95.2	-94.4	-93.6	-92.8	-92.0	-91.2	-90.4	-89.6

Figure 6: Model Output Tables for 0% - 5% Damages from SRM with L2C Controls Baseline

The first table again shows the results when SRM has damages amounting to 0% of gross world output. In this case, researching SRM provides greater benefits than the no research alternative for all probabilities of achieving optimal controls as well as for all probabilities of deploying the technology after research. As in the earlier scenario where L2C controls were assumed in the no research alternative, there are differences in total warming when SRM is used and when it is not. For instance, when there is no research, damages from the 2°C temperature increase will occur. When SRM is deployed in this case however, it counters all warming and avoids these temperature-related damages.

When compared to the immediately preceding scenario, the difference seen at 0% damages from SRM use can be explained by the change in the emissions control scenario when research is not pursued. When looking next at the case where SRM use causes 1% damages, there is still a sizable region where research provides greater benefits than the no research alternative. At smaller probabilities of achieving L2C controls, the cheaper no controls policy is employed. For these small probabilities, the reduced damages from employing a no controls emissions control scenario compensate for the 1% damages caused by deploying the technology and net to greater benefits when research is pursued than when not, as shown in the green shaded area of the output table.

For the remaining cases of 2% - 5% damages from SRM, research only provides equivalent or larger benefits than no research when the technology has zero probability of use. When that holds, any probability of achieving L2C controls will create equivalent or larger benefits for research than no research because the no controls emissions policy has smaller damages than L2C. The moral hazard risk compensation argument is then again irrelevant to the quality of our decision at these damage levels. It should also be noted that when SRM is researched and not used in this scenario, smaller probabilities of achieving L2C emissions controls have greater benefits than certain (p=1) L2C emissions controls.

5.2 SRM Offset Level 2°C

5.2.1 Optimal Emissions Controls

Sample Tree for 2% SRM Damages

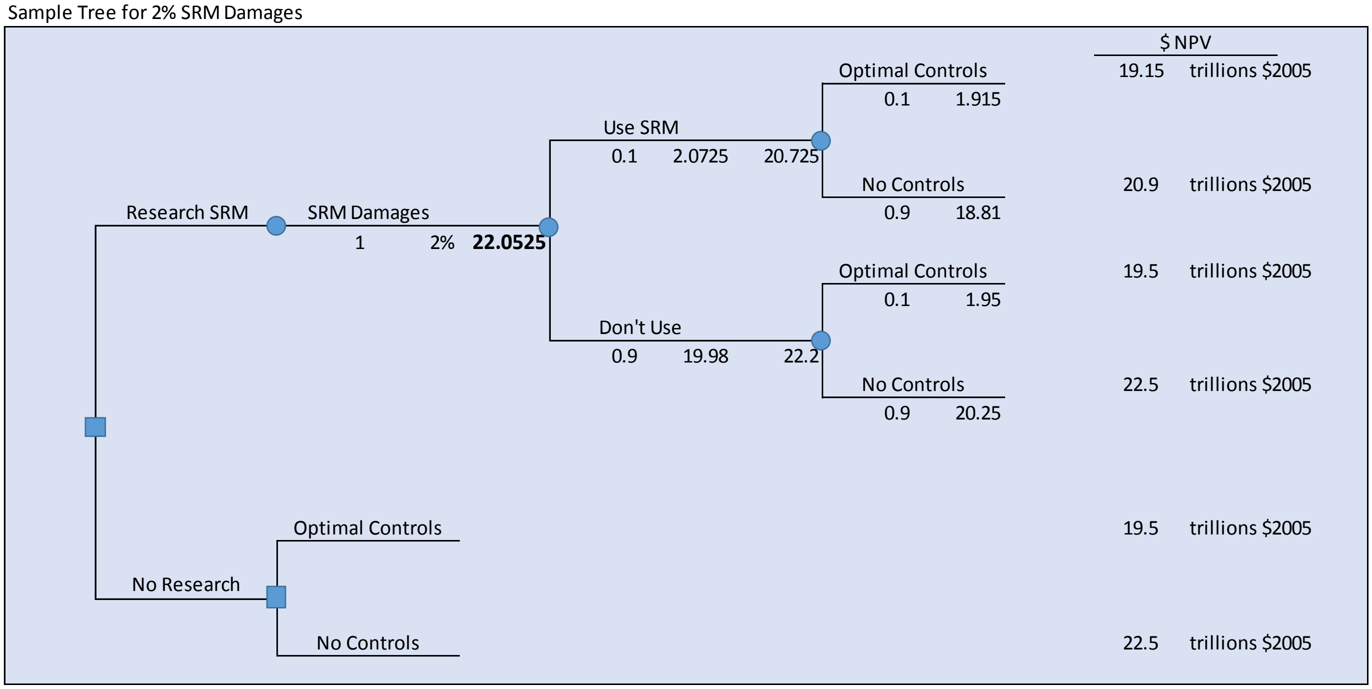

Figure 7: Sample Tree for 2% SRM Damages when SRM Level is Offset 2°C with Optimal Emissions Controls

In the following two scenarios SRM is used to limit temperature change to 2°C. The trade-off set for emissions controls after research is conducted includes no controls and optimal controls. That implies that, when compared with the complete offset scenarios in the previous section, the following scenarios use less SRM. This is an interesting comparison because both supporters and critics of research make note of the levels of SRM used, whether in discussion of optimal use or in fear of inappropriate use in conflict.

The decision tree above shows the sample case when research for SRM reveals that there will be certain damages of 2% of gross world output. As in the previous sections, the model assumes there is uncertainty as to whether SRM will be deployed and uncertainty about the emissions control strategy adopted after pursuing research.

The following pages show the decision to pursue research when SRM will be used at levels to limit temperature change to 2°C. As before, each page of outputs is a comparison to an individual emissions control strategy with SRM damages ranging from 0% to 5% of gross world output. Moving from left to right across each table, the probability of optimal controls increases by 0.1 per cell. Moving from the top to the bottom of the table, the probability of deploying SRM increases by 0.1 per cell. The green cells represent a good decision to research SRM and the red cells a bad decision.

Emission Control Level No Controls

This page shows our decision when SRM is used to offset warming to 2°C compared to a no controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research includes optimal controls and no controls.

0% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ / Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.7	3.0
0.1	0.8	1.0	1.3	1.5	1.8	2.1	2.3	2.6	2.8	3.1	3.3
0.2	1.5	1.8	2.0	2.2	2.4	2.6	2.8	3.0	3.3	3.5	3.7
0.3	2.3	2.5	2.7	2.8	3.0	3.2	3.3	3.5	3.7	3.9	4.0
0.4	3.1	3.2	3.3	3.5	3.6	3.7	3.9	4.0	4.1	4.3	4.4
0.5	3.9	3.9	4.0	4.1	4.2	4.3	4.4	4.5	4.6	4.6	4.7
0.6	4.6	4.7	4.7	4.8	4.8	4.9	4.9	4.9	5.0	5.0	5.1
0.7	5.4	5.4	5.4	5.4	5.4	5.4	5.4	5.4	5.4	5.4	5.4
0.8	6.2	6.1	6.1	6.0	6.0	6.0	5.9	5.9	5.8	5.8	5.8
0.9	6.9	6.9	6.8	6.7	6.6	6.5	6.4	6.4	6.3	6.2	6.1
1.0	7.7	7.6	7.5	7.3	7.2	7.1	7.0	6.8	6.7	6.6	6.5

1% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ / Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.7	3.0
0.1	0.5	0.7	1.0	1.3	1.6	1.8	2.1	2.4	2.6	2.9	3.2
0.2	0.9	1.2	1.4	1.7	1.9	2.2	2.4	2.6	2.9	3.1	3.4
0.3	1.4	1.6	1.8	2.1	2.3	2.5	2.7	2.9	3.1	3.3	3.6
0.4	1.9	2.1	2.3	2.4	2.6	2.8	3.0	3.2	3.4	3.6	3.7
0.5	2.4	2.5	2.7	2.8	3.0	3.1	3.3	3.5	3.6	3.8	3.9
0.6	2.8	3.0	3.1	3.2	3.3	3.5	3.6	3.7	3.9	4.0	4.1
0.7	3.3	3.4	3.5	3.6	3.7	3.8	3.9	4.0	4.1	4.2	4.3
0.8	3.8	3.8	3.9	4.0	4.1	4.1	4.2	4.3	4.3	4.4	4.5
0.9	4.2	4.3	4.3	4.4	4.4	4.5	4.5	4.5	4.6	4.6	4.7
1.0	4.7	4.7	4.7	4.8	4.8	4.8	4.8	4.8	4.8	4.8	4.9

2% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ / Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.7	3.0
0.1	0.2	0.4	0.7	1.0	1.3	1.6	1.9	2.2	2.5	2.7	3.0
0.2	0.3	0.6	0.9	1.1	1.4	1.7	2.0	2.2	2.5	2.8	3.1
0.3	0.5	0.7	1.0	1.3	1.5	1.8	2.1	2.3	2.6	2.8	3.1
0.4	0.6	0.9	1.1	1.4	1.6	1.9	2.1	2.4	2.6	2.9	3.1
0.5	0.8	1.0	1.3	1.5	1.8	2.0	2.2	2.5	2.7	2.9	3.2
0.6	1.0	1.2	1.4	1.6	1.9	2.1	2.3	2.5	2.8	3.0	3.2
0.7	1.1	1.3	1.5	1.8	2.0	2.2	2.4	2.6	2.8	3.0	3.2
0.8	1.3	1.5	1.7	1.9	2.1	2.3	2.5	2.7	2.9	3.1	3.3
0.9	1.4	1.6	1.8	2.0	2.2	2.4	2.6	2.8	2.9	3.1	3.3
1.0	1.6	1.8	2.0	2.1	2.3	2.5	2.7	2.8	3.0	3.2	3.4

3% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ / Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.7	3.0
0.1	-0.2	0.1	0.4	0.8	1.1	1.4	1.7	2.0	2.3	2.6	2.9
0.2	-0.3	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.8
0.3	-0.5	-0.2	0.1	0.5	0.8	1.1	1.4	1.7	2.0	2.3	2.6
0.4	-0.6	-0.3	0.0	0.3	0.6	0.9	1.2	1.6	1.9	2.2	2.5
0.5	-0.8	-0.5	-0.2	0.2	0.5	0.8	1.1	1.4	1.7	2.1	2.4
0.6	-1.0	-0.6	-0.3	0.0	0.3	0.6	1.0	1.3	1.6	1.9	2.3
0.7	-1.1	-0.8	-0.5	-0.1	0.2	0.5	0.8	1.2	1.5	1.8	2.1
0.8	-1.3	-1.0	-0.6	-0.3	0.0	0.4	0.7	1.0	1.3	1.7	2.0
0.9	-1.4	-1.1	-0.8	-0.4	-0.1	0.2	0.5	0.9	1.2	1.5	1.9
1.0	-1.6	-1.3	-0.9	-0.6	-0.3	0.1	0.4	0.7	1.1	1.4	1.8

4% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ / Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.7	3.0
0.1	-0.5	-0.2	0.2	0.5	0.8	1.1	1.4	1.8	2.1	2.4	2.7
0.2	-1.0	-0.6	-0.3	0.1	0.4	0.7	1.1	1.4	1.8	2.1	2.4
0.3	-1.4	-1.1	-0.7	-0.4	0.0	0.4	0.7	1.1	1.4	1.8	2.1
0.4	-1.9	-1.5	-1.2	-0.8	-0.4	0.0	0.3	0.7	1.1	1.5	1.9
0.5	-2.4	-2.0	-1.6	-1.2	-0.8	-0.4	0.0	0.4	0.8	1.2	1.6
0.6	-2.9	-2.5	-2.0	-1.6	-1.2	-0.8	-0.4	0.0	0.5	0.9	1.3
0.7	-3.4	-2.9	-2.5	-2.1	-1.6	-1.2	-0.7	-0.3	0.1	0.6	1.0
0.8	-3.8	-3.4	-2.9	-2.5	-2.0	-1.6	-1.1	-0.6	-0.2	0.3	0.7
0.9	-4.3	-3.8	-3.4	-2.9	-2.4	-1.9	-1.5	-1.0	-0.5	0.0	0.4
1.0	-4.8	-4.3	-3.8	-3.3	-2.8	-2.3	-1.8	-1.3	-0.8	-0.3	0.1

5% SRM Damages vs No SRM + No Controls

P (Use SRM) ↓ / Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	0.3	0.6	0.9	1.2	1.5	1.8	2.1	2.4	2.7	3.0
0.1	-0.8	-0.5	-0.1	0.2	0.5	0.9	1.2	1.5	1.9	2.2	2.6
0.2	-1.6	-1.2	-0.9	-0.5	-0.1	0.3	0.6	1.0	1.4	1.7	2.1
0.3	-2.4	-2.0	-1.6	-1.2	-0.8	-0.4	0.0	0.4	0.9	1.3	1.7
0.4	-3.2	-2.8	-2.3	-1.9	-1.4	-1.0	-0.5	-0.1	0.3	0.8	1.2
0.5	-4.0	-3.5	-3.0	-2.6	-2.1	-1.6	-1.1	-0.7	-0.2	0.3	0.8
0.6	-4.8	-4.3	-3.8	-3.3	-2.7	-2.2	-1.7	-1.2	-0.7	-0.2	0.3
0.7	-5.6	-5.1	-4.5	-4.0	-3.4	-2.9	-2.3	-1.8	-1.2	-0.7	-0.1
0.8	-6.4	-5.8	-5.2	-4.6	-4.1	-3.5	-2.9	-2.3	-1.7	-1.1	-0.6
0.9	-7.2	-6.6	-6.0	-5.3	-4.7	-4.1	-3.5	-2.9	-2.2	-1.6	-1.0
1.0	-8.0	-7.3	-6.7	-6.0	-5.4	-4.7	-4.1	-3.4	-2.8	-2.1	-1.5

Figure 8: Model Output Tables for 0% - 5% Damages from SRM with No Controls Baseline

At a first glance, the outputs in this scenario are significantly different than those in the earlier scenarios where SRM was used to completely offset any warming. Two important factors which account for these differences are the trade-off set for emissions controls after research is completed and the assumed emissions control policy when there is no research. For the cases with 0% - 2% damages from SRM use, researching SRM provides greater benefits than the no research alternative for all probabilities of achieving optimal controls as well as for all probabilities of deploying the technology after research.

Even for the cases with 3% - 5% damages from SRM use there are very large areas where research provides greater benefits. When optimal controls are achieved, there is very little SRM use needed to limit temperature increase to 2°C. When there are low probabilities of achieving optimal controls and high probabilities of using the technology, more SRM is needed to compensate for the larger warming in the no controls situation. In this case, there are fewer benefits from research than in the no research alternative, as shown in the bottom left corner of the output tables for 3% - 5% damages from SRM.

Emission Control Level Optimal Controls

This page shows our decision when SRM is used to offset warming to 2°C compared to an optimal controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research again includes optimal controls and no controls.

0% SRM Damages vs No SRM + Optimal Controls

P (Use SRM) ↓ \ Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-2.2	-2.0	-1.7	-1.5	-1.2	-0.9	-0.7	-0.4	-0.2	0.1	0.3
0.2	-1.5	-1.2	-1.0	-0.8	-0.6	-0.4	-0.2	0.0	0.3	0.5	0.7
0.3	-0.7	-0.5	-0.3	-0.1	0.0	0.2	0.3	0.5	0.7	0.9	1.0
0.4	0.1	0.2	0.3	0.5	0.6	0.7	0.9	1.0	1.1	1.3	1.4
0.5	0.9	0.9	1.0	1.1	1.2	1.3	1.4	1.5	1.6	1.6	1.7
0.6	1.6	1.7	1.7	1.8	1.8	1.9	1.9	1.9	2.0	2.0	2.1
0.7	2.4	2.4	2.4	2.4	2.4	2.4	2.4	2.4	2.4	2.4	2.4
0.8	3.2	3.1	3.1	3.0	3.0	3.0	2.9	2.9	2.8	2.8	2.8
0.9	3.9	3.9	3.8	3.7	3.6	3.5	3.4	3.4	3.3	3.2	3.1
1.0	4.7	4.6	4.5	4.3	4.2	4.1	4.0	3.8	3.7	3.6	3.5

1% SRM Damages vs No SRM + Optimal Controls

P (Use SRM) ↓ \ Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-2.5	-2.3	-2.0	-1.7	-1.4	-1.2	-0.9	-0.6	-0.4	-0.1	0.2
0.2	-2.1	-1.8	-1.6	-1.3	-1.1	-0.8	-0.6	-0.4	-0.1	0.1	0.4
0.3	-1.6	-1.4	-1.2	-0.9	-0.7	-0.5	-0.3	-0.1	0.1	0.3	0.6
0.4	-1.1	-0.9	-0.7	-0.6	-0.4	-0.2	0.0	0.2	0.4	0.6	0.7
0.5	-0.6	-0.5	-0.3	-0.2	0.0	0.1	0.3	0.5	0.6	0.8	0.9
0.6	-0.2	0.0	0.1	0.2	0.3	0.5	0.6	0.7	0.9	1.0	1.1
0.7	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	1.1	1.2	1.3
0.8	0.8	0.8	0.9	1.0	1.1	1.1	1.2	1.3	1.3	1.4	1.5
0.9	1.2	1.3	1.3	1.4	1.4	1.5	1.5	1.5	1.6	1.6	1.7
1.0	1.7	1.7	1.7	1.8	1.8	1.8	1.8	1.8	1.8	1.8	1.9

2% SRM Damages vs No SRM + Optimal Controls

P (Use SRM) ↓ \ Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-2.8	-2.6	-2.3	-2.0	-1.7	-1.4	-1.1	-0.8	-0.5	-0.3	0.0
0.2	-2.7	-2.4	-2.1	-1.9	-1.6	-1.3	-1.0	-0.8	-0.5	-0.2	0.1
0.3	-2.5	-2.3	-2.0	-1.7	-1.5	-1.2	-0.9	-0.7	-0.4	-0.2	0.1
0.4	-2.4	-2.1	-1.9	-1.6	-1.4	-1.1	-0.9	-0.6	-0.4	-0.1	0.1
0.5	-2.2	-2.0	-1.7	-1.5	-1.3	-1.0	-0.8	-0.5	-0.3	-0.1	0.2
0.6	-2.0	-1.8	-1.6	-1.4	-1.1	-0.9	-0.7	-0.5	-0.2	0.0	0.2
0.7	-1.9	-1.7	-1.5	-1.2	-1.0	-0.8	-0.6	-0.4	-0.2	0.0	0.2
0.8	-1.7	-1.5	-1.3	-1.1	-0.9	-0.7	-0.5	-0.3	-0.1	0.1	0.3
0.9	-1.6	-1.4	-1.2	-1.0	-0.8	-0.6	-0.4	-0.2	-0.1	0.1	0.3
1.0	-1.4	-1.2	-1.1	-0.9	-0.7	-0.5	-0.4	-0.2	0.0	0.2	0.3

3% SRM Damages vs No SRM + Optimal Controls

P (Use SRM) ↓ \ Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-3.2	-2.9	-2.6	-2.2	-1.9	-1.6	-1.3	-1.0	-0.7	-0.4	-0.1
0.2	-3.3	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3
0.3	-3.5	-3.2	-2.9	-2.5	-2.2	-1.9	-1.6	-1.3	-1.0	-0.7	-0.4
0.4	-3.6	-3.3	-3.0	-2.7	-2.4	-2.1	-1.8	-1.4	-1.1	-0.8	-0.5
0.5	-3.8	-3.5	-3.2	-2.8	-2.5	-2.2	-1.9	-1.6	-1.3	-0.9	-0.6
0.6	-4.0	-3.6	-3.3	-3.0	-2.7	-2.4	-2.0	-1.7	-1.4	-1.1	-0.8
0.7	-4.1	-3.8	-3.5	-3.1	-2.8	-2.5	-2.2	-1.8	-1.5	-1.2	-0.9
0.8	-4.3	-4.0	-3.6	-3.3	-3.0	-2.6	-2.3	-2.0	-1.7	-1.3	-1.0
0.9	-4.4	-4.1	-3.8	-3.4	-3.1	-2.8	-2.5	-2.1	-1.8	-1.5	-1.1
1.0	-4.6	-4.3	-3.9	-3.6	-3.3	-2.9	-2.6	-2.3	-1.9	-1.6	-1.3

4% SRM Damages vs No SRM + Optimal Controls

P (Use SRM) ↓ \ Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-3.5	-3.2	-2.8	-2.5	-2.2	-1.9	-1.6	-1.2	-0.9	-0.6	-0.3
0.2	-4.0	-3.6	-3.3	-2.9	-2.6	-2.3	-1.9	-1.6	-1.2	-0.9	-0.6
0.3	-4.4	-4.1	-3.7	-3.4	-3.0	-2.6	-2.3	-1.9	-1.6	-1.2	-0.9
0.4	-4.9	-4.5	-4.2	-3.8	-3.4	-3.0	-2.7	-2.3	-1.9	-1.5	-1.1
0.5	-5.4	-5.0	-4.6	-4.2	-3.8	-3.4	-3.0	-2.6	-2.2	-1.8	-1.4
0.6	-5.9	-5.5	-5.0	-4.6	-4.2	-3.8	-3.4	-3.0	-2.5	-2.1	-1.7
0.7	-6.4	-5.9	-5.5	-5.1	-4.6	-4.2	-3.7	-3.3	-2.9	-2.4	-2.0
0.8	-6.8	-6.4	-5.9	-5.5	-5.0	-4.6	-4.1	-3.6	-3.2	-2.7	-2.3
0.9	-7.3	-6.8	-6.4	-5.9	-5.4	-4.9	-4.5	-4.0	-3.5	-3.0	-2.6
1.0	-7.8	-7.3	-6.8	-6.3	-5.8	-5.3	-4.8	-4.3	-3.8	-3.3	-2.9

5% SRM Damages vs No SRM + Optimal Controls

P (Use SRM) ↓ \ Q (OC) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	-3.0	-2.7	-2.4	-2.1	-1.8	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.1	-3.8	-3.5	-3.1	-2.8	-2.5	-2.1	-1.8	-1.5	-1.1	-0.8	-0.4
0.2	-4.6	-4.2	-3.9	-3.5	-3.1	-2.7	-2.4	-2.0	-1.6	-1.3	-0.9
0.3	-5.4	-5.0	-4.6	-4.2	-3.8	-3.4	-3.0	-2.6	-2.1	-1.7	-1.3
0.4	-6.2	-5.8	-5.3	-4.9	-4.4	-4.0	-3.5	-3.1	-2.7	-2.2	-1.8
0.5	-7.0	-6.5	-6.0	-5.6	-5.1	-4.6	-4.1	-3.7	-3.2	-2.7	-2.2
0.6	-7.8	-7.3	-6.8	-6.3	-5.7	-5.2	-4.7	-4.2	-3.7	-3.2	-2.7
0.7	-8.6	-8.1	-7.5	-7.0	-6.4	-5.9	-5.3	-4.8	-4.2	-3.7	-3.1
0.8	-9.4	-8.8	-8.2	-7.6	-7.1	-6.5	-5.9	-5.3	-4.7	-4.1	-3.6
0.9	-10.2	-9.6	-9.0	-8.3	-7.7	-7.1	-6.5	-5.9	-5.2	-4.6	-4.0
1.0	-11.0	-10.3	-9.7	-9.0	-8.4	-7.7	-7.1	-6.4	-5.8	-5.1	-4.5

Figure 9: Model Output Tables for 0% - 5% Damages from SRM with Optimal Controls Baseline

Looking at this scenario with optimal controls as the baseline when there is no research compared with the immediately preceding scenario where no controls are used as the baseline when no research is performed, it becomes clear that the baseline used for comparison makes a substantial difference in the decision to pursue research. In the output tables for 0% and 1% damages from SRM use, there is a sizable area where the probability of using the technology is low and the probability of achieving optimal controls is also low. In these cases the damages from mitigating too little are not compensated by the benefits from using SRM, and pursuing research provides fewer benefits than the no research alternative.

When there are 2% damages for SRM use, there is no penalty for research if the probability of achieving optimal controls is one. There are even a few cases where the probability of achieving optimal controls is slightly lower but the certain or near-certain use of the technology compensates for the increased damages incurred by less than optimal mitigation. For the cases where there are 3% - 5% damages for SRM use, research only provides equivalent benefits to no research when there is zero probability the technology will be used and achieving optimal emissions controls is certain.

5.2.2 Limit 2°C Emissions Controls

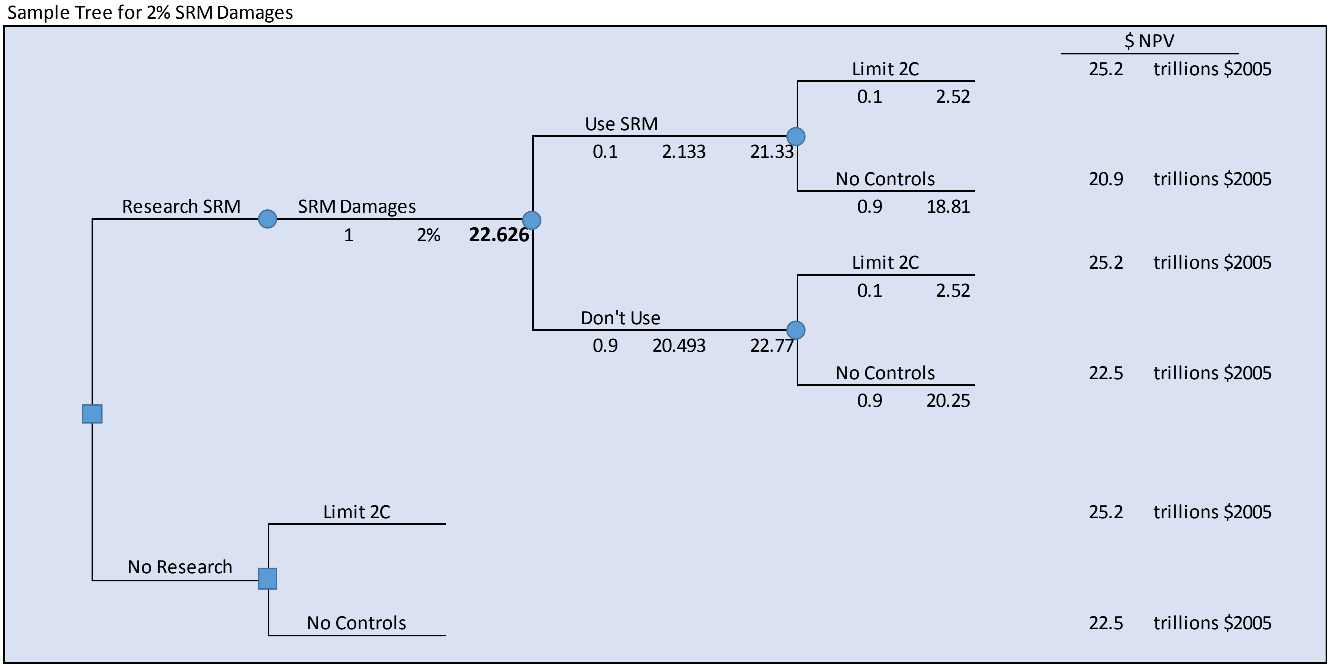

Figure 10: Sample Tree for 2% SRM Damages when SRM Level is Offset 2°C with Limit 2°C Controls

In the following two scenarios SRM is again used to limit temperature change to 2°C. The trade-off set for emissions controls after research is conducted includes no controls and optimal controls. Figure 10 above shows the decision tree for the sample case when research for SRM reveals that there will be certain damages of 2% of gross world output. As in the previous sections, the model assumes there is uncertainty as to whether SRM will be deployed and uncertainty about the emissions control strategy adopted after pursuing research.

Again, each page of outputs is a comparison to an individual emissions control strategy with SRM damages ranging from 0% to 5% of gross world output. Moving from left to right across each table, the probability of optimal controls increases by 0.1 per cell.

Moving from the top to the bottom of the table, the probability of deploying SRM increases by 0.1 per cell. The green cells represent a good decision to research SRM and the red cells a bad decision.

Emission Control Level No Controls

This page shows our decision when SRM is used to offset warming to 2°C compared to a no controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research includes L2C controls and no controls.

0% SRM Damages vs No SRM + No Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	0.8	0.4	0.1	-0.3	-0.6	-1.0	-1.3	-1.7	-2.0	-2.4	-2.7
0.2	1.5	1.1	0.7	0.3	-0.2	-0.6	-1.0	-1.4	-1.9	-2.3	-2.7
0.3	2.3	1.8	1.3	0.8	0.3	-0.2	-0.7	-1.2	-1.7	-2.2	-2.7
0.4	3.1	2.5	1.9	1.3	0.8	0.2	-0.4	-1.0	-1.5	-2.1	-2.7
0.5	3.9	3.2	2.5	1.9	1.2	0.6	-0.1	-0.7	-1.4	-2.0	-2.7
0.6	4.6	3.9	3.2	2.4	1.7	1.0	0.2	-0.5	-1.2	-2.0	-2.7
0.7	5.4	4.6	3.8	3.0	2.2	1.3	0.5	-0.3	-1.1	-1.9	-2.7
0.8	6.2	5.3	4.4	3.5	2.6	1.7	0.8	0.0	-0.9	-1.8	-2.7
0.9	6.9	6.0	5.0	4.0	3.1	2.1	1.2	0.2	-0.8	-1.7	-2.7
1.0	7.7	6.7	5.6	4.6	3.5	2.5	1.5	0.4	-0.6	-1.7	-2.7

1% SRM Damages vs No SRM + No Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	0.5	0.2	-0.2	-0.5	-0.8	-1.1	-1.4	-1.7	-2.1	-2.4	-2.7
0.2	0.9	0.6	0.2	-0.2	-0.5	-0.9	-1.2	-1.6	-2.0	-2.3	-2.7
0.3	1.4	1.0	0.6	0.2	-0.2	-0.6	-1.1	-1.5	-1.9	-2.3	-2.7
0.4	1.9	1.4	1.0	0.5	0.1	-0.4	-0.9	-1.3	-1.8	-2.2	-2.7
0.5	2.4	1.8	1.3	0.8	0.3	-0.2	-0.7	-1.2	-1.7	-2.2	-2.7
0.6	2.8	2.3	1.7	1.2	0.6	0.1	-0.5	-1.0	-1.6	-2.1	-2.7
0.7	3.3	2.7	2.1	1.5	0.9	0.3	-0.3	-0.9	-1.5	-2.1	-2.7
0.8	3.8	3.1	2.5	1.8	1.2	0.5	-0.1	-0.8	-1.4	-2.1	-2.7
0.9	4.2	3.5	2.9	2.2	1.5	0.8	0.1	-0.6	-1.3	-2.0	-2.7
1.0	4.7	4.0	3.2	2.5	1.7	1.0	0.3	-0.5	-1.2	-2.0	-2.7

2% SRM Damages vs No SRM + No Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	0.2	-0.1	-0.4	-0.7	-1.0	-1.3	-1.6	-1.8	-2.1	-2.4	-2.7
0.2	0.3	0.0	-0.3	-0.6	-0.9	-1.2	-1.5	-1.8	-2.1	-2.4	-2.7
0.3	0.5	0.2	-0.2	-0.5	-0.8	-1.1	-1.4	-1.7	-2.1	-2.4	-2.7
0.4	0.6	0.3	0.0	-0.4	-0.7	-1.0	-1.4	-1.7	-2.0	-2.4	-2.7
0.5	0.8	0.4	0.1	-0.3	-0.6	-1.0	-1.3	-1.7	-2.0	-2.4	-2.7
0.6	1.0	0.6	0.2	-0.1	-0.5	-0.9	-1.2	-1.6	-2.0	-2.3	-2.7
0.7	1.1	0.7	0.4	0.0	-0.4	-0.8	-1.2	-1.6	-1.9	-2.3	-2.7
0.8	1.3	0.9	0.5	0.1	-0.3	-0.7	-1.1	-1.5	-1.9	-2.3	-2.7
0.9	1.4	1.0	0.6	0.2	-0.2	-0.6	-1.0	-1.5	-1.9	-2.3	-2.7
1.0	1.6	1.2	0.7	0.3	-0.1	-0.6	-1.0	-1.4	-1.8	-2.3	-2.7

3% SRM Damages vs No SRM + No Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-0.2	-0.4	-0.7	-0.9	-1.2	-1.4	-1.7	-1.9	-2.2	-2.4	-2.7
0.2	-0.3	-0.6	-0.8	-1.0	-1.3	-1.5	-1.7	-2.0	-2.2	-2.5	-2.7
0.3	-0.5	-0.7	-0.9	-1.1	-1.4	-1.6	-1.8	-2.0	-2.3	-2.5	-2.7
0.4	-0.6	-0.8	-1.1	-1.3	-1.5	-1.7	-1.9	-2.1	-2.3	-2.5	-2.7
0.5	-0.8	-1.0	-1.2	-1.4	-1.6	-1.8	-1.9	-2.1	-2.3	-2.5	-2.7
0.6	-1.0	-1.1	-1.3	-1.5	-1.7	-1.8	-2.0	-2.2	-2.4	-2.5	-2.7
0.7	-1.1	-1.3	-1.4	-1.6	-1.8	-1.9	-2.1	-2.2	-2.4	-2.5	-2.7
0.8	-1.3	-1.4	-1.6	-1.7	-1.8	-2.0	-2.1	-2.3	-2.4	-2.6	-2.7
0.9	-1.4	-1.6	-1.7	-1.8	-1.9	-2.1	-2.2	-2.3	-2.4	-2.6	-2.7
1.0	-1.6	-1.7	-1.8	-1.9	-2.0	-2.2	-2.3	-2.4	-2.5	-2.6	-2.7

4% SRM Damages vs No SRM + No Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-0.5	-0.7	-0.9	-1.1	-1.4	-1.6	-1.8	-2.0	-2.3	-2.5	-2.7
0.2	-1.0	-1.1	-1.3	-1.5	-1.7	-1.8	-2.0	-2.2	-2.4	-2.5	-2.7
0.3	-1.4	-1.6	-1.7	-1.8	-1.9	-2.1	-2.2	-2.3	-2.4	-2.6	-2.7
0.4	-1.9	-2.0	-2.1	-2.2	-2.2	-2.3	-2.4	-2.5	-2.5	-2.6	-2.7
0.5	-2.4	-2.4	-2.5	-2.5	-2.5	-2.6	-2.6	-2.6	-2.6	-2.7	-2.7
0.6	-2.9	-2.9	-2.8	-2.8	-2.8	-2.8	-2.8	-2.8	-2.7	-2.7	-2.7
0.7	-3.4	-3.3	-3.2	-3.2	-3.1	-3.0	-3.0	-2.9	-2.8	-2.8	-2.7
0.8	-3.8	-3.7	-3.6	-3.5	-3.4	-3.3	-3.2	-3.0	-2.9	-2.8	-2.7
0.9	-4.3	-4.2	-4.0	-3.8	-3.7	-3.5	-3.3	-3.2	-3.0	-2.9	-2.7
1.0	-4.8	-4.6	-4.4	-4.2	-4.0	-3.8	-3.5	-3.3	-3.1	-2.9	-2.7

5% SRM Damages vs No SRM + No Controls

P (Use SRM) \ Q (L2C)	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.0	-0.3	-0.5	-0.8	-1.1	-1.4	-1.6	-1.9	-2.2	-2.4	-2.7
0.1	-0.8	-1.0	-1.2	-1.4	-1.6	-1.8	-1.9	-2.1	-2.3	-2.5	-2.7
0.2	-1.6	-1.7	-1.8	-1.9	-2.0	-2.2	-2.3	-2.4	-2.5	-2.6	-2.7
0.3	-2.4	-2.4	-2.5	-2.5	-2.5	-2.6	-2.6	-2.6	-2.6	-2.7	-2.7
0.4	-3.2	-3.2	-3.1	-3.1	-3.0	-3.0	-2.9	-2.9	-2.8	-2.8	-2.7
0.5	-4.0	-3.9	-3.7	-3.6	-3.5	-3.4	-3.2	-3.1	-3.0	-2.8	-2.7
0.6	-4.8	-4.6	-4.4	-4.2	-4.0	-3.8	-3.5	-3.3	-3.1	-2.9	-2.7
0.7	-5.6	-5.3	-5.0	-4.7	-4.4	-4.2	-3.9	-3.6	-3.3	-3.0	-2.7
0.8	-6.4	-6.0	-5.7	-5.3	-4.9	-4.6	-4.2	-3.8	-3.4	-3.1	-2.7
0.9	-7.2	-6.8	-6.3	-5.9	-5.4	-5.0	-4.5	-4.1	-3.6	-3.2	-2.7
1.0	-8.0	-7.5	-6.9	-6.4	-5.9	-5.4	-4.8	-4.3	-3.8	-3.2	-2.7

Figure 11: Model Output Tables for 0% - 5% Damages from SRM with No Controls Baseline

This scenario is somewhat similar to the scenario in section 5.1.2 in that the trade-off set for emissions controls after pursing research includes L2C controls and no controls, and this is compared to a policy of no controls in the no research alternative. As before the target emissions policy, L2C, provides for higher damages than the no controls policy because the mitigation efforts necessary for L2C are inefficient. When high probabilities of achieving L2C controls occur with low probabilities of using the technology, the research option provides fewer benefits than the no research alternative, as seen in the output tables for 0% - 2% damages for SRM use.

After this point the only case where research provides equivalent or greater benefits is when the probabilities of using the technology and achieving L2C controls are both zero. This is again a somewhat odd comparison because the target emissions control policy has higher damages than the no controls policy, making the probability of the moral hazard risk compensation objection more beneficial to society. Though odd it is a very relevant comparison because, as mentioned in the introduction to the report, the most widely accepted international goal for preventing climate change damages is currently set at limiting temperature increases to 2°C.

Emission Control Level Limit 2°C

This page shows our decision when SRM is used to offset warming to 2°C compared to a L2C controls emissions scenario when there is no SRM research conducted. The trade-off set for emissions controls after research includes L2C controls and no controls.

0% SRM Damages vs No SRM + L2C Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	3.5	3.1	2.8	2.4	2.1	1.7	1.4	1.0	0.7	0.3	0.0
0.2	4.2	3.8	3.4	3.0	2.5	2.1	1.7	1.3	0.8	0.4	0.0
0.3	5.0	4.5	4.0	3.5	3.0	2.5	2.0	1.5	1.0	0.5	0.0
0.4	5.8	5.2	4.6	4.0	3.5	2.9	2.3	1.7	1.2	0.6	0.0
0.5	6.6	5.9	5.2	4.6	3.9	3.3	2.6	2.0	1.3	0.7	0.0
0.6	7.3	6.6	5.9	5.1	4.4	3.7	2.9	2.2	1.5	0.7	0.0
0.7	8.1	7.3	6.5	5.7	4.9	4.0	3.2	2.4	1.6	0.8	0.0
0.8	8.9	8.0	7.1	6.2	5.3	4.4	3.5	2.7	1.8	0.9	0.0
0.9	9.6	8.7	7.7	6.7	5.8	4.8	3.9	2.9	1.9	1.0	0.0
1.0	10.4	9.4	8.3	7.3	6.2	5.2	4.2	3.1	2.1	1.0	0.0

1% SRM Damages vs No SRM + L2C Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	3.2	2.9	2.5	2.2	1.9	1.6	1.3	1.0	0.6	0.3	0.0
0.2	3.6	3.3	2.9	2.5	2.2	1.8	1.5	1.1	0.7	0.4	0.0
0.3	4.1	3.7	3.3	2.9	2.5	2.1	1.6	1.2	0.8	0.4	0.0
0.4	4.6	4.1	3.7	3.2	2.8	2.3	1.8	1.4	0.9	0.5	0.0
0.5	5.1	4.5	4.0	3.5	3.0	2.5	2.0	1.5	1.0	0.5	0.0
0.6	5.5	5.0	4.4	3.9	3.3	2.8	2.2	1.7	1.1	0.6	0.0
0.7	6.0	5.4	4.8	4.2	3.6	3.0	2.4	1.8	1.2	0.6	0.0
0.8	6.5	5.8	5.2	4.5	3.9	3.2	2.6	1.9	1.3	0.6	0.0
0.9	6.9	6.2	5.6	4.9	4.2	3.5	2.8	2.1	1.4	0.7	0.0
1.0	7.4	6.7	5.9	5.2	4.4	3.7	3.0	2.2	1.5	0.7	0.0

2% SRM Damages vs No SRM + L2C Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	2.9	2.6	2.3	2.0	1.7	1.4	1.1	0.9	0.6	0.3	0.0
0.2	3.0	2.7	2.4	2.1	1.8	1.5	1.2	0.9	0.6	0.3	0.0
0.3	3.2	2.9	2.5	2.2	1.9	1.6	1.3	1.0	0.6	0.3	0.0
0.4	3.3	3.0	2.7	2.3	2.0	1.7	1.3	1.0	0.7	0.3	0.0
0.5	3.5	3.2	2.8	2.5	2.1	1.8	1.4	1.1	0.7	0.4	0.0
0.6	3.7	3.3	2.9	2.6	2.2	1.8	1.5	1.1	0.7	0.4	0.0
0.7	3.8	3.4	3.1	2.7	2.3	1.9	1.5	1.1	0.8	0.4	0.0
0.8	4.0	3.6	3.2	2.8	2.4	2.0	1.6	1.2	0.8	0.4	0.0
0.9	4.1	3.7	3.3	2.9	2.5	2.1	1.7	1.2	0.8	0.4	0.0
1.0	4.3	3.9	3.4	3.0	2.6	2.2	1.7	1.3	0.9	0.4	0.0

3% SRM Damages vs No SRM + L2C Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	2.5	2.3	2.0	1.8	1.5	1.3	1.0	0.8	0.5	0.3	0.0
0.2	2.4	2.1	1.9	1.7	1.4	1.2	1.0	0.7	0.5	0.2	0.0
0.3	2.2	2.0	1.8	1.6	1.3	1.1	0.9	0.7	0.4	0.2	0.0
0.4	2.1	1.9	1.6	1.4	1.2	1.0	0.8	0.6	0.4	0.2	0.0
0.5	1.9	1.7	1.5	1.3	1.1	1.0	0.8	0.6	0.4	0.2	0.0
0.6	1.7	1.6	1.4	1.2	1.0	0.9	0.7	0.5	0.3	0.2	0.0
0.7	1.6	1.4	1.3	1.1	0.9	0.8	0.6	0.5	0.3	0.2	0.0
0.8	1.4	1.3	1.1	1.0	0.9	0.7	0.6	0.4	0.3	0.1	0.0
0.9	1.3	1.1	1.0	0.9	0.8	0.6	0.5	0.4	0.3	0.1	0.0
1.0	1.1	1.0	0.9	0.8	0.7	0.6	0.4	0.3	0.2	0.1	0.0

4% SRM Damages vs No SRM + L2C Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	2.2	2.0	1.8	1.6	1.3	1.1	0.9	0.7	0.4	0.2	0.0
0.2	1.7	1.6	1.4	1.2	1.0	0.9	0.7	0.5	0.3	0.2	0.0
0.3	1.3	1.1	1.0	0.9	0.8	0.6	0.5	0.4	0.3	0.1	0.0
0.4	0.8	0.7	0.6	0.5	0.5	0.4	0.3	0.2	0.2	0.1	0.0
0.5	0.3	0.3	0.2	0.2	0.2	0.2	0.1	0.1	0.1	0.0	0.0
0.6	-0.2	-0.2	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	0.0	0.0	0.0
0.7	-0.7	-0.6	-0.5	-0.5	-0.4	-0.3	-0.3	-0.2	-0.1	-0.1	0.0
0.8	-1.1	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.3	-0.2	-0.1	0.0
0.9	-1.6	-1.5	-1.3	-1.1	-1.0	-0.8	-0.6	-0.5	-0.3	-0.2	0.0
1.0	-2.1	-1.9	-1.7	-1.5	-1.3	-1.1	-0.8	-0.6	-0.4	-0.2	0.0

5% SRM Damages vs No SRM + L2C Controls

P (Use SRM) ↓ \ Q (L2C) →	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	2.7	2.4	2.2	1.9	1.6	1.4	1.1	0.8	0.5	0.3	0.0
0.1	1.9	1.7	1.5	1.3	1.1	0.9	0.8	0.6	0.4	0.2	0.0
0.2	1.1	1.0	0.9	0.8	0.7	0.6	0.4	0.3	0.2	0.1	0.0
0.3	0.3	0.3	0.2	0.2	0.2	0.2	0.1	0.1	0.1	0.0	0.0
0.4	-0.5	-0.4	-0.4	-0.3	-0.3	-0.3	-0.2	-0.1	-0.1	-0.1	0.0
0.5	-1.3	-1.2	-1.0	-0.9	-0.8	-0.6	-0.5	-0.4	-0.3	-0.1	0.0
0.6	-2.1	-1.9	-1.7	-1.5	-1.3	-1.1	-0.8	-0.6	-0.4	-0.2	0.0
0.7	-2.9	-2.6	-2.3	-2.0	-1.7	-1.5	-1.2	-0.9	-0.6	-0.3	0.0
0.8	-3.7	-3.3	-3.0	-2.6	-2.2	-1.9	-1.5	-1.1	-0.7	-0.4	0.0
0.9	-4.5	-4.1	-3.6	-3.2	-2.7	-2.3	-1.8	-1.4	-0.9	-0.4	0.0
1.0	-5.3	-4.8	-4.2	-3.7	-3.2	-2.7	-2.1	-1.6	-1.1	-0.5	0.0

Figure 12: Model Output Tables for 0% - 5% Damages from SRM with L2C Controls Baseline

Once more, comparing this scenario with the immediately preceding scenario, it is clear that the baseline used for comparison makes a substantial difference in the decision to pursue research. This situation is somewhat special compared to other scenarios examined thus far. The SRM level is set to limit temperature change to 2°C, and the target emissions control policy is also set to limit temperature change to 2°C. That means when L2C controls are achieved, SRM is never used. This can be seen in the rightmost column of the output tables, where the difference in benefits between the two alternatives is zero. Just as in the previous scenario, the target emissions control policy has higher damages than the no controls policy.

For the cases with 0% - 3% damages from SRM use, there are equivalent or greater benefits for research than for the no research alternative for all probabilities of achieving L2C controls and all probabilities of using the technology. In the cases with 4% and 5% damages from SRM use, the damages from higher probabilities of SRM use are large enough that they outweigh the benefits of lower mitigation costs under the no controls policy.

5.3 Discussion of the Results

When comparing each of the previous emissions control and SRM use combinations, two points become apparent. First, as proposed in earlier sections, the amount of SRM used makes a considerable difference in the potential benefits of pursuing research. Second, beliefs about the emissions control strategy society to which society would adhere in the event research is not pursued make an even larger difference in the potential benefits of pursuing research.

There are key weaknesses to the model which limit its applicability as a policymaking tool. First, the model is based on an existing model with substantial uncertainties. Nordhaus [8] identifies the largest uncertainties in the DICE model as: the growth in global output, future development of energy systems, the pace of technological change to replace fossil fuel systems, climate sensitivity to increasing levels of greenhouse gases, and economic and ecological responses to climate change. Another weakness is the categorical nature of the trade-off policies when adjusting the probability of achieving the target emissions policy. That means rather than saying society achieved 80% of its emissions mitigation target policy, this model only allows society to completely meet the requirements of the target policy or meet none of the requirements while adjusting which category society falls in. An analysis of this nature was beyond the scope of this report, but would make for an interesting study as it more closely resembles the behavior of society than the binary policy trade-off presented here.

Chapter 6: Conclusion

This report has highlighted the most relevant considerations affecting society's decision to pursue climate engineering research. Sulfate aerosols warrant attention for their unique role as a risk management tool with the ability to cheaply and rapidly cool the Earth in the case of rapid warming. Substantial uncertainties in the costs and benefits associated with the technology and its use have led to ethical and moral objections to both researching and deploying the technology.

While the sensitivity model in this report analyzes the influence of two main objections to research on the decision to pursue research, that climate engineering research would lessen society's mitigation efforts and would necessarily lead to deployment of the technology, the outputs of the model show that the selection of the best alternative requires information about preferences and beliefs about society. For instance, preference for an optimal controls policy over a policy to limit temperature change to 2°C as a target emissions policy would generate a steeper penalty for failing to achieve the emissions target, thereby reducing the potential benefits of SRM research. Similarly, a belief that in the absence of SRM research society would achieve optimal controls reduces the apparent benefits of SRM research when compared with a belief that in the absence of SRM research society would achieve L2C controls.

There are limitations to extrapolating from the current situation, but there is also insight to be gained by looking at current progress in achieving internationally set emissions targets. In this case that sets the preference to a policy which limits temperature change to 2°C as in sections 5.1.2 and 5.2.2. If SRM would be used to completely offset warming, as in section 5.1.2, the decision to research SRM changes with the introduction of even very small damages from its use. If instead SRM would be used to limit warming to 2°C, the same temperature target as that of the emissions control policy, researching SRM proves itself a good decision across all damages when the baseline control policy in the absence of research is also L2C. In the case where the baseline in the absence of research is no controls, SRM research remains a good decision for most cases with small damages caused by SRM use. Neither of these scenarios provide reason to eliminate researching climate engineering as an alternative.

Ethical considerations for justice were not considered in the decision sensitivity model, though they may still influence the decision to pursue research. Svoboda, Keller, Goes, and Tuana [22] provided the most in-depth arguments for distributive, intergenerational, and procedural justice in the case of climate engineering. The framework provided in the study evaluated whether an alternative would meet the requirements of various theories of justice, though the authors chose to focus only on climate engineering. The authors raise interesting questions about the role of ethics in climate change policy, but beyond being an interesting thought experiment the framework established in their study provides little insight to a policymaker, because no alternative for managing climate change risk could pass the stringent requirements in the framework. Certainly policymakers should consider issues of justice, but the justice-based objections presented are also not sufficient to eliminate researching climate engineering as an alternative.

As a final consideration, many studies conclude with recommendations for climate engineering governance and suggestions for international cooperation. These recommendations include a broad range of ideas, including compensation schemes for those facing potential damages, intellectual property rights agreements, and public participation in the decision-making process [11, 28]. The recommendations appear to be an attempt to influence the decision to research SRM, and for that reason, I will not make specific recommendations about issues of governance. The decision on how to govern the technology, though undeniably important, needs to remain a separate decision from the decision to pursue research.

References

1. Arrow, K.J., *Global Climate Change: A Challenge to Policy.* The Economists' Voice, 2007. **4**(3): p. 2-5.
2. Reynolds, J., *A critical examination of the climate engineering moral hazard and risk compensation concern.* The Anthropocene Review, 2014: p. 1-18.
3. Edenofer, O., et al., *IPCC, 2014: Summary for Policymakers*. 2014: Cambridge.
4. Brasseur, G.P. and C. Granier, *Mitigation, Adaptation or Climate Engineering?* Theoretical Inquiries in Law, 2013. **14**(1): p. 1.
5. EPA, U. *A Student's Guide to Global Climate Change*. 2014 [cited 2015 1/5/2015]; Available from: http://www.epa.gov/climatestudents/solutions/technologies/index.html.
6. Arvesen, A., R.M. Bright, and E.G. Hertwich, *Considering only first-order effects? how simplifications lead to unrealistic technology optimism in climate change mitigation.* Energy Policy, 2011. **39**(11): p. 7448.
7. Barrett, S., *Proposal for a New Climate Change Treaty System.* The Economists' Voice, 2007. **4**(3): p. 6-6.
8. Nordhaus, W.D., *A question of balance: weighing the options on global warming policies*. 2008, New Haven: Yale University Press.
9. Keith, D.W., *A case for climate engineering*. 2013, Cambridge, Mass: The MIT Press.
10. Schelling, T.C., *The economic diplomacy of geoengineering.* Climatic Change, 1996. **33**(3): p. 303-307.
11. Rayner, S., et al., *The Oxford Principles.* Climatic Change, 2013. **121**(3): p. 499-512.
12. Barrett, S., *The incredible economics of geoengineering.* Environmental & resource economics, 2008. **39**(1): p. 45-54.
13. Vaughan, N. and T. Lenton, *A review of climate geoengineering proposals.* Climatic Change, 2011. **109**(3-4): p. 745-790.
14. Burns, W., *Into the Great Wide Open? The Potential Promise & Perils of Climate Geoengineering*, in *AESS Webinar Series*. 2015.

15. Crutzen, P.J., *Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma?* Climatic Change, 2006. **77**(3): p. 211-220.

16. Nibleus, K. and R. Lundin, *Climate change and mitigation.* Ambio, 2010. **39**(1): p. 11-17.

17. Forster, P. *Not enough time for geoengineering to work?* Bulletin of the Atomic Scientists, 2015.

18. Bickel, J.E., *Climate engineering and climate tipping-point scenarios.* Environment Systems & Decisions, 2013. **33**(1): p. 152.

19. Salter, S., G. Sortino, and J. Latham, *Sea-going hardware for the cloud albedo method of reversing global warming.* Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 2008. **366**(1882): p. 3989-4006.

20. Szulczewski, M.L., et al., *Lifetime of carbon capture and storage as a climate-change mitigation technology.* Proceedings of the National Academy of Sciences, 2012. **109**(14): p. 5185-5189.

21. Dechezleprêtre, A., et al., *Invention and transfer of climate change-mitigation technologies: a global analysis.* Review of environmental economics and policy, 2011. **5**(1): p. 109-130.

22. Svoboda, T.K., Klaus; Goes, Marlos; Tuana, Nancy, *Sulfate Aerosol Geoengineering: The Question of Justice.* Public Affairs Quarterly, 2011. **25**(3): p. 157-174.

23. Gardiner, S.M., *A perfect moral storm: the ethical tragedy of climate change*. 2011, New York: Oxford University Press.

24. Agency, U.S.E.P., *Updating Ozone Calculations and Emissions Profiles for use in the Atmospheric and Health Effects Framework Model*. 2015.

25. Lomborg, B.n., *Smart solutions to climate change: comparing costs and benefits*. 2010, New York: Cambridge University Press.

26. Goes, M., N. Tuana, and K. Keller, *The economics (or lack thereof) of aerosol geoengineering.* Climatic Change, 2011. **109**(3): p. 719-744.

27. Morrow, D.R., *Ethical aspects of the mitigation obstruction argument against climate engineering research.* Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 2014. **372**(2031): p. 20140062.

28. Betz, G., *The case for climate engineering research: an analysis of the "arm the future" argument.* Climatic Change, 2012. **111**(2): p. 473-485.

29. Lomborg, B.n., *Global problems, smart solutions: costs and benefits*. 2013, Cambridge: Cambridge University Press.

30. Wong, P.-H., *Distributive Justice, Geoengineering and Risks.* Climate Geoengineering Governance Working Papers, 2014.

31. Bickel, J.E. and S. Agrawal, *Reexamining the economics of aerosol geoengineering.* Climatic Change, 2013. **119**(3): p. 993-1006.

32. Collier, P., *The plundered planet: why we must, and how we can, manage nature for global prosperity*. 2010, Oxford: Oxford University Press.

33. Eickmann, L.G.K. and R.A.N. Morisetti. *Climate Change Impacts as a Catalyst for Conflict*. 2015. LBJ School of Public Affairs.

34. Kilic, C., M. Buckholtz, Editor. 2015.