



INVESTMENT IN PEOPLE AND IDEAS

CHILDREN AND FAMILIES
EDUCATION AND THE ARTS
ENERGY AND ENVIRONMENT
HEALTH AND HEALTH CARE
INFRASTRUCTURE AND
TRANSPORTATION
INTERNATIONAL AFFAIRS
LAW AND BUSINESS
NATIONAL SECURITY
POPULATION AND AGING
PUBLIC SAFETY
SCIENCE AND TECHNOLOGY
TERRORISM AND
HOMELAND SECURITY

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis.

This electronic document was made available from www.rand.org as a public service of the RAND Corporation.

Skip all front matter: [Jump to Page 1](#) ▼

Support RAND

[Browse Reports & Bookstore](#)

[Make a charitable contribution](#)

For More Information

Visit RAND at www.rand.org

Explore the [RAND Investment in People and Ideas Program](#)

View [document details](#)


Limited Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law as indicated in a notice appearing later in this work. This electronic representation of RAND intellectual property is provided for non-commercial use only. Unauthorized posting of RAND electronic documents to a non-RAND website is prohibited. RAND electronic documents are protected under copyright law. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents for commercial use. For information on reprint and linking permissions, please see [RAND Permissions](#).

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011	
4. TITLE AND SUBTITLE Governing Geoengineering Research: A Political and Technical Vulnerability Analysis of Potential Near-Term Options				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) RAND Corporation, 1776 Main Street, P.O. Box 2138, Santa Monica, CA, 90407-2138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 95	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

This product is part of the RAND Corporation technical report series. Reports may include research findings on a specific topic that is limited in scope; present discussions of the methodology employed in research; provide literature reviews, survey instruments, modeling exercises, guidelines for practitioners and research professionals, and supporting documentation; or deliver preliminary findings. All RAND reports undergo rigorous peer review to ensure that they meet high standards for research quality and objectivity.

TECHNICAL
REPORT



Governing Geoengineering Research

A Political and Technical
Vulnerability Analysis of Potential
Near-Term Options

Robert J. Lempert, Don Prosnitz



INVESTMENT IN PEOPLE AND IDEAS

This report is a product of the RAND Corporation's continuing program of self-initiated independent research. Support for such research is provided, in part, by donors and by the independent research and development provisions of RAND's contracts for the operation of its U.S. Department of Defense federally funded research and development centers. The research was conducted jointly within the RAND Environment, Energy, and Economic Development Program of RAND Infrastructure, Safety, and Environment and the Acquisition and Technology Policy Center of the RAND National Defense Research Institute.

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

© Copyright 2011 RAND Corporation

Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Copies may not be duplicated for commercial purposes. Unauthorized posting of RAND documents to a non-RAND website is prohibited. RAND documents are protected under copyright law. For information on reprint and linking permissions, please visit the RAND permissions page (<http://www.rand.org/publications/permissions.html>).

Published 2011 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213-2665
RAND URL: <http://www.rand.org>
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Preface

Geoengineering, the large-scale, persistent, intentional altering of the globe's climate, is increasingly mentioned as a potential response to climate change. But evaluating the risks associated with any policy toward geoengineering confronts deep uncertainties concerning not only the desirability of deploying such systems but also the consequences of conducting research and large-scale experiments on climate modification. This report demonstrates a decision framework for conducting a risk analysis under such conditions of deep uncertainty. It also provides an initial evaluation of alternative near-term policies the U.S. government might pursue regarding the governance of geoengineering research. The analysis utilizes a vulnerability-and-response-option analysis decision framework, based on robust decisionmaking methods and utilizing a simple Bayesian belief net simulation, to compare the risks of alternative policy choices. The robust decisionmaking concepts employed in this report have been widely used by the RAND Corporation to address many climate-related and other policy questions but heretofore have not been applied to questions of geoengineering.

This report should be of interest to policymakers, scholars, and researchers interested in the governance of geoengineering research, development, eventual deployment, and, especially, the security implications of geoengineering.

This report is a product of the RAND Corporation's continuing program of self-initiated independent research. Support for such research is provided, in part, by donors and by the independent research and development provisions of RAND's contracts for the operation of its U.S. Department of Defense federally funded research and development centers. The research was conducted jointly within the RAND Environment, Energy, and Economic Development Program (EEED) of RAND Infrastructure, Safety, and Environment (ISE) and the Acquisition and Technology Policy Center of the RAND National Defense Research Institute.

The RAND National Defense Research Institute is a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, the Unified Combatant Commands, the Navy, the Marine Corps, the defense agencies, and the defense Intelligence Community. For more information on the Acquisition and Technology Policy Center, see <http://www.rand.org/nsrd/about/atp.html> or contact the director (contact information is provided on the web page).

The mission of ISE is to improve the development, operation, use, and protection of society's essential physical assets and natural resources and to enhance the related social assets of safety and security of individuals in transit and in their workplaces and communities. The EEED research portfolio addresses environmental quality and regulation, energy resources and systems, water resources and systems, climate, natural hazards and disasters, and economic development—both domestically and internationally. EEED research is conducted for govern-

ment, foundations, and the private sector. Information about the Environment, Energy, and Economic Development Program is available online (<http://www.rand.org/ise/environ>). Inquiries about EEED projects should be sent to the following address:

Keith Crane, Director
Environment, Energy, and Economic Development Program, ISE
RAND Corporation
1200 South Hayes Street
Arlington, VA 22202-5050
703-413-1100, x5520
Keith_Crane@rand.org

Questions or comments about this report should be sent to the project leader, Robert Lempert (Robert_Lempert@rand.org).

Contents

Preface	iii
Figures	vii
Tables	ix
Summary	xi
Acknowledgments	xiii
Abbreviations	xv
 CHAPTER ONE	
Introduction	1
 CHAPTER TWO	
Governance and Security Challenges	5
Geoengineering Deployments and Their Risks	6
Managing Geoengineering with Research Norms	9
 CHAPTER THREE	
A Vulnerability-and-Response-Option Analysis Framework for a Risk Assessment of Geoengineering Governance	15
Key Factors Considered in the Analysis	17
Levers	18
Uncertainties	19
Measures	20
Relationships	22
Caveats	27
Evaluating Policies Across Multiple Scenarios	28
 CHAPTER FOUR	
Identifying the Vulnerabilities of Alternative U.S. Geoengineering Governance Policies	35
 CHAPTER FIVE	
Conclusions and Next Steps	39
 APPENDIXES	
A. Types of Geoengineering Technologies	43
B. Relationships	47
C. Conditional Probability Tables	51
 References	 73

Figures

2.1.	Five Phases of a Geoengineering Program	9
3.1.	Simulation Nodes for Strong Norms Policy	23
3.2.	Simulation Nodes for Ban	24
3.3.	Simulation Nodes for No Norms Policy	25
3.4.	Relative Regret (%) of Three Policies Over 16 Scenarios.....	31

Tables

3.1.	Factors Considered in the Analysis: Exogenous Uncertainties (X), Policy Levers (L), Measures (M), and Relationships (R)	18
3.2.	Range of Cases Considered for Each Uncertain Parameter	28
3.3.	Performance for Three Policies for the Scenario Shown in Figure 3.1.....	30
3.4.	Relative Regret for Three Policies for the Scenario Shown in Figure 3.1.....	30
4.1.	Summary of Vulnerabilities of Alternative Policies	37
B.1.	Relationships Among Nodes in Bayesian Network.....	47
C.1.	Emission-Reduction Influence Factors	52
C.2.	Conditional Probability Table: DeploymentAgreement	53
C.3.	Conditional Probability Table: EmissionsReductions	53
C.4.	Conditional Probability Table: TestResults.....	55
C.5.	Conditional Probability Table: OthersTest	56
C.6.	Conditional Probability Table: OthersDeploy	56
C.7.	Conditional Probability Table: USTests	59
C.8.	Conditional Probability Table: DeploymentResults.....	60
C.9.	Conditional Probability Table: ClimateImpactsUS.....	63
C.10.	Conditional Probability Table: CollateralDamage	64
C.11.	Conditional Probability Table: Conflict	65
C.12.	Conditional Probability Table: ClimateImpactsOthers.....	67
C.13.	Conditional Probability Table: USDeploys.....	68

Summary

Geoengineering—the deliberate altering of the earth’s climate—represents a risky and, for many, a frightening proposition. But the concept has attracted increasing interest in recent years because of its potential ability to significantly transform the portfolio of options for limiting the magnitude of future climate change. In contrast to most approaches for reducing greenhouse gas emissions, some geoengineering approaches could prove fast acting and inexpensive and could be deployed by one or a small number of nations without global cooperation. These characteristics present significant challenges for risk management, national security, and international governance that have only just begun to be seriously considered.

This report provides an initial examination and comparison of the risks associated with alternative international approaches the United States might pursue to governing solar radiation management (SRM) geoengineering. To handle the extensive, wide-ranging uncertainties, we employ a vulnerability-and-response-option analysis decision framework. Specifically, we identify scenarios in which alternative U.S. policies toward geoengineering governance might fail to meet their goals and suggest how alternative policies might reduce those vulnerabilities. The report implements this approach using a simple simulation model to conduct the first steps of a robust decisionmaking (RDM) analysis. The analysis identifies some of the risks of three commonly debated near-term approaches to managing geoengineering research: establishing strong norms for research, banning research entirely, or leaving research unregulated.

This report aims to serve three purposes. First, it demonstrates the potential ability for a risk analysis based on a vulnerability-and-response-option analysis framework to inform the debate on geoengineering. Second, it helps define the steps needed to conduct a full RDM analysis to address such governance issues. Third, it provides some intriguing, if only suggestive, policy results.

This analysis compared three alternative policies the U.S. government might pursue regarding near-term geoengineering governance. The report focuses on SRM technologies because these technologies offer the full range of characteristics that make geoengineering both so alluring and dangerous: possibly fast acting, potentially relatively inexpensive, and likely to cause global consequences from even unilateral action.

Under Strong Norms, the U.S. government would encourage the establishment of international norms to govern geoengineering research. Under Ban, the United States would promote a prohibition on any geoengineering research. Under No Norms, the United States would actively discourage the formation of norms governing research. A comparison of the performance of these three strategies across many plausible future states of the world suggests that, if U.S. policymakers believe that some type of SRM technology is possible, they ought to prefer the Strong Norms policy to No Norms or Ban. Under such conditions, this option outper-

forms the alternatives because it increases the likelihood of a successful deployment in those cases in which geoengineering proves useful. It also reduces the likelihood of failed deployments by nations struggling to respond to serious climate impacts.

If U.S. policymakers believe that no SRM geoengineering technology is likely to succeed, they might prefer the Ban or No Norms policy to Strong Norms. The Ban policy appears the better of the two if policymakers believe that climate change is highly unlikely to prove catastrophic. Under such conditions, this option reduces the risks of overconfidence—deploying a geoengineering system that passes its tests but fails in practice. This option also increases the likelihood of reaching an international agreement to reduce greenhouse gas emissions.

In contrast, U.S. policymakers might prefer No Norms to Strong Norms or Ban if they believe that SRM technologies are unlikely to work but that climate change could prove catastrophic. Under such conditions and a Ban policy, some other nation might defy the prohibition, test, and then deploy a SRM system that subsequently fails. The absence of research norms might lead to uncoordinated tests by several nations, undermining the ability to learn from any test. Thus, in this case, the absence of research norms might prove more effective than the Ban at preventing unsuccessful geoengineering deployments under dire circumstances.

Many caveats attend these initial results. The analysis considers only a small set of the options available to the U.S. government. The report focuses only on the decisions of national governments and does not explicitly consider the choices of private firms and other nongovernmental actors that might influence the evolution of geoengineering policies. A more-complete RDM analysis with an enhanced simulation model would likely suggest additional vulnerabilities beyond those identified here and likely identify ways to ameliorate at least some of them. However, this report does demonstrate an approach to risk analysis under conditions of deep uncertainty that, in an expanded form, could help U.S. policymakers develop and evaluate robust policies toward geoengineering governance. The study also offers some initial insights about the future conditions under which alternative approaches for governing geoengineering research might not perform as expected, provides some initial suggestions regarding the trade-offs among such strategies, and describes next steps that could result in a more-complete assessment of the trade-offs among alternative near-term policies for managing the risk and opportunities of geoengineering.

Acknowledgments

We are grateful for Charles Meade's assistance in conducting part of the research reported here. We gratefully acknowledge our reviewers James Bonomo and Jason Blackstock, whose thoughtful comments and criticisms helped us significantly improve this document. Lisa Bernard provided much-appreciated editorial assistance. We gratefully acknowledge the generous support of the RAND Corporation for conducting this work.

Abbreviations

ABP	assumption-based planning
BBN	Bayesian belief network
CDR	carbon dioxide removal
CO ₂	carbon dioxide
ppm	parts per million
R&D	research and development
RDM	robust decisionmaking
SRM	solar radiation management
SWCE	short-wave climate engineering

Introduction

Geoengineering—the deliberate altering of the earth’s climate—represents a risky and, for many, frightening proposition. But the concept has attracted increasing interest in recent years because of its potential ability to significantly transform the portfolio of options for limiting the magnitude of future climate change. In particular, some studies estimate that some geoengineering approaches could cost orders of magnitude less than reducing greenhouse gas emissions and could show environmental benefits much sooner—in months to years, rather than in decades (Barrett, 2008; Victor, 2008; Royal Society, 2009). Furthermore, it has been suggested that one or a few nations could deploy geoengineering systems without cooperation from the rest of the world (Victor et al., 2009). These features make geoengineering potentially attractive given the emerging scientific evidence that even a small rise in global temperatures might pose significant dangers and given the potential difficulties of achieving large global reductions of greenhouse gases over the next few decades.

But geoengineering also presents significant risks. Many approaches aimed at reducing global temperatures could have potentially severe adverse side effects. Some impacts, such as exacerbating droughts or eroding the ozone layer, are already well recognized by scientists even if those scientists disagree about the seriousness of the potential problems (Keith, 2000; Bala, Duffy, and Taylor, 2008; Caldeira and Wood, 2008). Other adverse impacts could remain complete surprises until a geoengineering system is deployed. In contrast to many investments aimed at reducing emissions, which might be difficult to implement but that, once in place, are relatively easy to sustain,¹ many geoengineering approaches would require extensive efforts to maintain over centuries or more. Relaxing such efforts could lead to rapid and catastrophic climate impacts as the global climate quickly warms, precluding social and environmental adaptation (Goes, Keller, and Tuana, 2010). In addition, because some nations could implement geoengineering unilaterally, the approach raises an entirely different set of political and security concerns from those raised by efforts to limit global emissions. With geoengineering, one or several nations could quite plausibly improve the climate for themselves while making it much worse for others (Victor, 2008).

These characteristics of potential geoengineering systems present significant challenges for risk management, national security, and international governance.² Over the past few years, an increasing number of articles and studies have begun to suggest that the best near-term

¹ Lempert, Popper, Resetar, and Hart (2002) report that, once made, many investments in emission-producing infrastructure can remain in place for many decades, despite changing economic and regulatory conditions.

² Recent discussions of the security and governance challenges of geoengineering include Victor (2008), Lin (2009), Barrett (2008), and Fleming (2007).

approach to control these risks requires national governments and scientific societies to establish research norms—that is, codes of conduct—for any research on geoengineering systems. For instance, a recent report by the Royal Society (2009, p. xii) recommends the following:

The Royal Society in collaboration with international scientific partners should develop a code of practice for geoengineering research and provide recommendations to the international scientific community for a voluntary research governance framework. . . . This should include . . . the establishment of a *de minimus* [sic] standard for the regulation of research.³

The Royal Society and other commentators (Cicerone, 2006; Victor, 2008; Lin, 2009) reject an outright ban on geoengineering research because, in some particularly dire climate-change scenarios, geoengineering might prove necessary and the lack of extensive preparatory research might render potentially attractive options unavailable or more likely to fail if nonetheless deployed. These commentators also stress the dangers of leaving geoengineering unregulated, emphasizing the need for some international rules governing research to prevent environmental damage from tests and to smooth the way for future international agreements governing any deployment of such systems.

Although there have been some initial attempts to evaluate the risks of individual geoengineering technologies (see, for instance, Royal Society, 2009; Boyd, 2008), few have attempted a systematic risk assessment and comparison of alternative approaches for governing geoengineering research, testing, and development.⁴ In large part, this derives from the complicated and uncertain string of events such research and development might initiate. Most geoengineering technologies exist today only as notional concepts sketched by scientists. The future impacts that any potential systems might have on those who deploy them and on others worldwide will, at best, remain deeply uncertain until the principles underlying proposed systems have been rigorously tested and the system engineering completed. More troubling, given the complexity of the climate system, is the fact that many geoengineering deployments will likely produce impossible-to-predict surprises, some relatively benign and others potentially significant. The potential for such unknown unknowns complicates any serious assessment of geoengineering governance. In addition, technical uncertainties are only a small set of the full range of uncertainties surrounding geoengineering. At least as important are political questions, such as the impacts a vigorous geoengineering program might have on international efforts to control greenhouse gas emissions, the willingness of various nations to deploy geoengineering in an attempt to protect themselves against climate-change impacts, and how other nations would respond if they perceived such deployments to be adversely affecting their interests. Few good precedents exist to provide any solid basis for analysts' judgments about how nations might react in such novel and potentially trying circumstances.

This report provides an initial attempt at a systematic risk assessment of near-term policies regarding geoengineering governance that the U.S. government might pursue. The analysis is clearly preliminary, but this report aims to serve three useful purposes. First, any attempt to

³ Also see the results of a recent expert meeting at Asilomar that attempted to craft guidelines for geoengineering research (Kintisch, 2010).

⁴ Some recent analyses—for instance, Bickel and Lane (2009)—still largely neglect any evaluation of geoengineering risks in their cost-benefit analyses. In contrast, other commentators focus almost entirely on the risks of geoengineering (Fleming, 2007).

assess the risks of geoengineering policies must confront the deep uncertainty that surrounds the issue. This report demonstrates the potential utility of an approach to risk analysis under such conditions of deep uncertainty, based on a vulnerability-and-response-option decision framework (Lempert, Nakicenovic, et al., 2004; Lempert, Groves, et al., 2006; Bryant and Lempert, 2010). Second, this report helps define the steps needed to conduct a more-complete analysis of U.S. geoengineering governance using this approach. Third, it provides some intriguing, if only suggestive, policy results. In particular, it illustrates the (often unstated) assumptions underlying arguments for alternative U.S. policy choices.

This study focuses on what the Royal Society calls solar radiation management (SRM) geoengineering technologies. The term geoengineering applies to a wide range of approaches. Some, which the Royal Society calls carbon dioxide (CO₂) removal (CDR), focus on extracting greenhouse gases from the atmosphere. As summarized in Appendix A, CDR technologies should prove relatively straightforward to manage because, as we discuss, they are slow acting, comparable in expense to many approaches for reducing emissions, and likely to cause any adverse effects only in the regions where they are deployed. In contrast, SRM technologies reflect a small amount of solar radiation back into space, thereby cooling the planet and offsetting some of the heat trapped by increasing concentrations of greenhouse gases. Furthermore, SRM will reduce the average global temperatures irrespective of the underlying cause of climate warming (although some regional changes could be detrimental). These technologies offer the full range of characteristics that make geoengineering both so alluring and so dangerous: possibly fast acting, potentially relatively inexpensive, and likely to cause global consequences from unilateral action. As described by Victor (2008, p. 323), SRM technologies could “turn the politics of climate protection upside down.”

This report employs a multiscenario vulnerability-and-response-option analysis to handle the extensive, wide-ranging uncertainties, including the potential for surprise (Lempert, Popper, and Bankes, 2002), that confront any analysis of geoengineering policy. In particular, the analysis aims to identify future conditions in which each of three commonly debated near-term approaches to managing geoengineering—establishing strong norms for research, banning research entirely, or leaving research unregulated—would fail to achieve U.S. policymakers’ goals. The resulting scenarios combine a range of key factors, including the future performance of geoengineering technologies, the seriousness of future climate impacts, the international community’s ability to jointly implement programs to reduce greenhouse gas emissions and regulate any geoengineering deployments, and the influence of geoengineering research programs on efforts to reduce greenhouse gas emissions. The semiquantitative vulnerability analysis offered by this report represents a partial implementation of a robust decision-making (RDM) analysis, an approach for quantitative policy analysis under conditions of deep uncertainty (Lempert, Popper, and Bankes, 2003).

This report’s preliminary vulnerability analysis suggests that the research-norm strategy proposed by the Royal Society and others might prove a reasonable near-term approach for managing SRM geoengineering technologies. But, for U.S. policymakers to pursue such a strategy, they should believe that some type of SRM approach is technically feasible. If, on the other hand, they believe that no SRM technology is likely to succeed, then they might alternatively prefer using a ban or even, under some assumptions about future climate conditions, a near-term strategy that blocks any international controls over research on such technologies. We must emphasize, however, that these are initial observations that require further study. For instance, a more-complete RDM analysis might identify ways in which policymakers could

modify a strategy promoting strong research norms in order to ameliorate some of the potential vulnerabilities described here.

The next chapter reviews the conditions that might lead one or more nations to deploy SRM technologies and the reasons that the United States might pursue a near-term norm-creation strategy to manage the risk of such technologies. The third chapter presents the vulnerability-and-response-option analysis methodology for examining alternative near-term U.S. government policies toward geoengineering governance. The fourth chapter presents the findings of such analysis. The final chapter summarizes the policy implications and suggests next steps that could result in a more-complete assessment of the trade-offs among alternative near-term policies for managing the risk and opportunities of geoengineering. Three appendixes survey geoengineering technologies and provide details of the simulation model used in this study.

Governance and Security Challenges

To understand the implications of alternative policies for governing geoengineering research, it is important to consider the circumstances that might lead a government to deploy a geoengineering system.

Geoengineering has gained interest because limiting the magnitude of climate change by reducing greenhouse gas emissions presents a difficult policy challenge. Atmospheric concentrations of greenhouse gases have been growing at about 0.4 percent annually for the past half-century, largely due to human combustion of fossil fuels. To date, global temperatures have risen by about 1 degree Celsius above preindustrial levels, driven in large part by this increase in greenhouse gas concentrations. Given the time lags in the climate system, primarily the large thermal inertia of the oceans and the longevity of carbon in the atmosphere, this warming would continue for decades or more even if all human emissions ceased immediately (Solomon et al., 2009). But, for more than two centuries, fossil-fuel use has powered economic growth and higher standards of living. At present, many developing countries use less than one-fifth of the energy per person than much richer developed nations. Thus, stabilizing and then reducing atmospheric concentrations of greenhouse gases in a growing global economy will require an unprecedented transformation of the world's energy, industrial, and transportation systems that, at best, will take decades.

Implementing policies to foster such a transformation has, not surprisingly, proved difficult. As of this writing, the European Union, some U.S. states, and a handful of national governments have legislated long-term goals for significant reductions of greenhouse gas emissions. But realizing these commitments remains a significant challenge even in these jurisdictions, and stabilizing concentrations will require global reductions, and a wide range of countries, including the United States and China, have yet to commit to such goals.

The consequences of failing to stabilize atmospheric concentrations of greenhouse gases also remain deeply uncertain. On one hand, even significant amounts of climate change might ultimately prove sufficiently benign that humans have little trouble adapting to any changes that occur. But the consequences of climate change could also prove severe (NRC, 2010a; Bernstein, Pachauri, and Reisinger, 2008). In particular, there is some evidence that the climate system might be approaching “tipping points” at which adverse changes begin to accelerate (NRC, 2002; Clark et al., 2008; Weitzman, 2009). Lenton et al. (2008) attempts to rank several climatic tipping points by the potential risk to society. More recently, Rockström et al. (2009) has attempted to define numeric planetary boundaries—points, such as atmospheric

CO₂ concentrations, beyond which society cannot afford to proceed.¹ In the worst cases, such tipping points could become unavoidable unless the most-ambitious goals for emission reductions over the next few decades are, in fact, realized (Hansen et al., 2008).

Geoengineering Deployments and Their Risks

Given this context, there are several situations that might lead one or more nations to deploy SRM geoengineering systems. In each case, the nations would be driven by the potentially low cost of such systems, the speed with which they could counter adverse effects of climate change, and the potential for such a system to prove effective without global participation in its deployment and operation.

In one situation that might lead to a geoengineering deployment, the world's nations might fail to sufficiently reduce emissions of greenhouse gases. As damage from climate change mounts, one or more nations might take matters into their own hands and deploy an SRM system in an attempt to halt impacts they found particularly detrimental. For instance, such a deployment might seek to save coastal areas at risk of submersion by halting the breakup of ice sheets and thermal expansion of the oceans, or might seek to break a long-term drought that engulfed some nation's agricultural lands. If deployed under these conditions, the geoengineering system would likely have to remain in operation for centuries to counteract the global warming resulting from the continued buildup of greenhouse gases.

In another situation that might lead to a geoengineering deployment, the world's nations might successfully hold future greenhouse gases to low levels. But even this achievement might prove insufficient to prevent the onset of catastrophic change if the climate system nears some tipping point. One or more nations, or many nations operating under an international agreement, might deploy an SRM system in an attempt to prevent catastrophic climate shifts. If deployed under these conditions, the geoengineering system would have to remain in operation for several decades while greenhouse gas concentrations peaked and finally began to decline.²

In another situation that might lead to a deployment, one or more nations might decide that geoengineering provides a much less expensive means for addressing any adverse impacts that climate change poses for their countries and deploy a system as an alternative to reducing greenhouse gases. If deployed under these conditions, the geoengineering system would likely have to remain in operation for centuries to counteract the continued buildup of greenhouse gases.

Despite any potential attractiveness, deployment and even research on such SRM systems in these or any other cases also pose some profound risks, which a small number of recent studies have only just begun to assess.

¹ Energy secretary Steven Chu (2009) has warned that many scientists fear that allowing atmospheric concentrations of CO₂ to exceed 450 parts per million (ppm) might lead to such an abrupt change:

The real danger with global warming will be the tipping point. As polar ice caps melt, the thaws could expose microbes, which would release carbon dioxide in quantities that would outstrip any reductions humans could make in their carbon-dioxide emissions.

² Deployment of CDR geoengineering technologies could reduce atmospheric greenhouse gas concentrations once emissions drop to sufficiently low levels.

The Royal Society (2009) and Boyd (2008) provide qualitative tabulations of the benefits and adverse impacts of various geoengineering technologies. The Royal Society evaluates the effectiveness, affordability, timeliness, and safety of such systems. Stratospheric aerosol injection technology, which would inject tiny particles into the stratosphere to reflect incoming solar radiation, achieves the highest ranking for the first three attributes but presents serious safety concerns, including adverse effects on regional hydrologic cycles, stratospheric ozone, high-altitude troposphere clouds, and biological productivity. (See Appendix A for a description of this and other SRM technologies.) Space-based methods, which use orbiting mirrors to reduce the sunlight reaching the earth, rank high on effectiveness but low on affordability and timeliness. The society assumes that it could take decades to deploy the necessary reflectors into space. Space-based methods rank medium on safety, reflecting possible impacts on hydrological cycles. The society ranks the cloud-albedo enhancement approach, which increases cloud cover by spraying fine particles into the air above the ocean, as medium on effectiveness, affordability, and timeliness, largely reflecting uncertainty about how such systems might perform. Testing might prove such system to be effective or not. The society gives cloud-albedo approaches a low ranking for safety, reflecting their potential impacts on weather patterns and ocean currents.

Boyd has similarly evaluated stratospheric aerosols and cloud-whitening SRM approaches (as well as the ocean fertilization, atmospheric carbon capture, and geochemical carbon-capture CDR approaches) using attributes similar to those used by the Royal Society, which he terms efficacy, affordability, safety, and rapidity. Boyd's rankings differ somewhat from the Royal Society's. For instance, in contrast to the society, Boyd gives cloud whitening, as well as stratospheric aerosols, the highest possible scores for their rapidity in slowing climate change, as well as their ability to implement an emergency stop. (Note that Boyd's rapidity attribute also includes the rate at which the geoengineering intervention could be shut down if adverse effects proved much larger than expected.) Boyd gives aerosols and cloud whitening similar scores for efficacy and ranks the latter higher for affordability. Similarly to the Royal Society, Boyd ranks cloud whitening as somewhat safer than stratospheric aerosols.

The differences between the Royal Society's and Boyd's rankings might, in large part, reflect the shortcomings of any semiquantitative ranking scheme that conflates judgments about the impacts of a technology, the probability of those impacts, and any uncertainty about the estimates of both the impacts and their probabilities. These differing estimates of impact and likelihood might be disentangled with more-quantitative risk assessments.

Goes, Keller, and Tuana (2010) has provided one of the few quantitative risk assessments of geoengineering systems. The authors use a simple integrated assessment model to compare the costs and benefits for sulfate aerosol geoengineering and greenhouse gas emission abatement as a function of uncertainties about the climate sensitivity, the damage from climate change, the costs of reducing greenhouse gas emissions, and the ability of future generations to maintain a geoengineering intervention once started. The first two uncertainties help characterize the potential impacts of climate change. The more serious and rapid the potential for such change, the more important it becomes to implement some combination of emission reductions and geoengineering. The third assumption characterizes the attractiveness of policies that rely on emission reductions. The fourth assumption characterizes some of the political uncertainties associated with deploying geoengineering. Goes and his coauthors find that two assumptions dominate any judgments about geoengineering's cost-effectiveness: the costs of any adverse side effects and the ability to maintain the system in place for centuries. An

SRM geoengineering system will appear cost-effective if decisionmakers are very certain—on the order of 90-percent confident—that it can be maintained for decades and that its adverse impacts will prove smaller than about 0.5 percent of gross world product. Otherwise, the study finds that the system's risks outweigh its benefits.

These quantitative results are consistent with the claims of many commentators that one of geoengineering's most-important risks could be political. One of the earliest and enduring criticisms is that any interest in the approach could undercut the political resolve needed to reduce emissions of greenhouse gases (Cicerone, Elliott, and Turco, 1992; Schneider, 1996; Schelling, 1996; Keith, 2000; Virgoe, 2008; Royal Society, 2009). This so-called moral hazard applies not only to deployment of such geoengineering systems but also to research programs that might make future deployments seem more politically and morally acceptable and thus more likely.³ Critics envision an insidious negative feedback in which initial investments in geoengineering research could decrease efforts toward reducing emissions, thereby increasing the incentives and need to pursue geoengineering. Any SRM geoengineering system confronts significant technical risks and potential adverse side effects and would fail to remedy at least some effects—most notably, ocean acidification (Doney et al., 2009)—of increasing CO₂ concentration. Thus, any policy that pursues geoengineering without any serious efforts at mitigation is sure to prove suboptimal. In addition, allowing greenhouse gas concentrations to grow over decades and centuries while offsetting the effects with increasing levels of geoengineering would place society and the planet in an increasingly precarious and unstable position. If the geoengineering system ever failed or were left unattended, global temperatures could rocket upward, generating climate change of unprecedented speed and scale.

Geoengineering systems also raise important international security concerns because, in some situations, they could increase the risk of conflict among nations. The effects of any SRM system could vary across different geographic regions and might ameliorate some consequences of climate change more than others. In each of the situations discussed above, one or a few nations could deploy SRM geoengineering to prevent some impact they find particularly onerous. But, in so doing, they might amplify impacts in some other region of the world—for instance, increasing droughts in other nations' agricultural lands or the frequency of extreme storms. Those affected, at best, might demand compensation and, at worst, seek to destroy or counteract the geoengineering system. The potential difficulty in attributing any drought or storms to the geoengineering system, as opposed to other climate change or natural variability, could exacerbate potential conflicts, making geoengineers and perceived victims less likely to even agree on what, if any, damage had been done.

Even if the disparate impacts of a geoengineering system were perfectly understood, the potential for prompt, low-cost, and unilateral control of the climate seriously exacerbates the question of whose hands control the global thermostat. The recognition that a few degrees of warming might greatly enhance agriculture in Siberia while turning California's agricultural heartland into a desert already complicates negotiations over emission reductions, even though any benefits remain uncertain and would likely prove fleeting if emissions continue to rise. But imagine if individual governments acquired the ability to tune the climate to their liking with sufficient speed that the resulting economic gains or losses became apparent in a few short

³ However, Polborn and Tintelnot (2009) model a case in which the potential for geoengineering can increase the chances for aggressive greenhouse gas reductions because, under some conditions, one nation can successfully use the threat to deploy geoengineering unilaterally to compel others to cooperate with emission-reduction efforts.

years. The pressures to tune the climate in ways that benefited one's own constituencies at the expense of others might prove strong.


Finally, there remains the possibility that one or more governments might use a geoengineering capability as a weapon, intentionally inflicting harm on some other nation.

Managing Geoengineering with Research Norms

These geoengineering risks have spawned much recent discussion, and some surprising consensus, on how geoengineering ought to be governed, at least in the near term. The near-term governance question is made easier because so little is known about the practical challenges and implications of geoengineering systems that the near-term questions all revolve around what type of laboratory research and field experiments ought to be conducted and under what type of oversight. The analysis presented here strongly suggests, however, that any evaluation of alternative options for near-term geoengineering governance should take into account the uncertain potential impacts such choices might have on future decisions about research, testing, and deployment.

Figure 2.1 shows five phases of any geoengineering program (Blackstock et al., 2009) and the level of environmental impact they might entail. A program might begin with noninvasive laboratory research. For instance, a stratospheric aerosol program might begin climate modeling studies of the potential impacts that such aerosols could have on hydrological cycles; engineering feasibility studies of alternative deployment options; and laboratory studies of the best size, chemical composition, and means to disperse the aerosols. None of these activities would have any environmental impacts.

Figure 2.1
Five Phases of a Geoengineering Program



SOURCE: Blackstock et al. (2009).

NOTE: SWCE = short-wave climate engineering.

RAND TR846-2.1

The next phase might include field experiments. A stratospheric aerosol program might begin with small-scale releases of aerosols to test their chemical and optical properties in actual atmospheric conditions. Some experiments have already been conducted (Izrael et al., 2009). Next might come larger releases maintained for weeks or months or even decades with attempts to monitor their global and regional impacts on the climate. Although such experiments would presumably be designed to have small environmental consequences and could be stopped quickly if needed, they would pose some level of environmental risk. There is great uncertainty as to whether experiments can be conducted at a large-enough scale to observe both the desired and undesired impacts—thereby significantly reducing the risks of a full-scale deployment without possessing unacceptable risk in and of themselves. Robock et al. (2010) argue that geoengineering cannot ever be effectively tested at subscale, thus implying that testing and deployment are, by necessity, one and the same.

The third phase might involve a scale-up to full deployment. Ideally, the process would be carefully monitored to ensure that all was unfolding as predicted. If problems were detected, deployment could be halted until they were corrected. This scale-up might unfold over the course of a few years. The fourth phase might involve the full geoengineering intervention, which might last for decades to centuries, depending on whether the geoengineering deployment were accompanied by significant reductions in greenhouse gas emissions. In the final phase, the deployment would be phased out.

Some commentators have advocated a ban on any geoengineering activities as the best approach to controlling the risks posed by geoengineering. Such a ban would certainly apply to any phase II field experiments and potentially to phase I laboratory work as well. Advocates see a ban as the best means for preventing interest in geoengineering from detracting from efforts to reduce greenhouse gas emissions. Some also favor a ban lest research lead to a system so potentially attractive, or bureaucratic momentum so strong, that deployment would become inevitable. But a ban also presents widely recognized risks. First, it could prove difficult to enforce. Would monitoring of research and sanctions be part of the ban? The only nations with geoengineering programs (thus violating the ban) might be those least likely to deploy any system cooperatively. Would countries considering their own best interests be willing to fall behind in a geoengineering race? In addition, although a ban might reduce the likelihood that a nation would deploy a geoengineering system in order to avoid making its own significant emission reductions, a ban might make more dangerous any deployment in response to the onset of an abrupt climate change or the failure of other nations to reduce emissions. In both of these latter cases, climate conditions might become so severe that some nations would feel compelled to attempt a geoengineering deployment. Without the knowledge gained by an extended period of testing, any system they deployed might prove far more likely to fail. In such situations, a geoengineering ban would increase the likelihood of a failed geoengineering deployment.

In place of an outright ban, some commentators have suggested pursuing a treaty or other international agreements as a means to regulate research and deployment of geoengineering systems (Barrett, 2008). But, although potentially useful in the long term, there are potentially strong arguments that attempts to negotiate such a treaty could prove counterproductive in the short term (Victor, 2008).

At present, several existing environmental treaties could, in principle, govern some geoengineering activities (Rayfuse, Lawrence, and Gjerde, 2008). The London Convention

and Protocol⁴ on ocean dumping was triggered to halt planned tests of ocean-fertilization approaches and subsequently modified to provide a framework for governing such tests (International Maritime Organization, 2007, 2008). The Geneva Convention on Transboundary Air Pollution⁵ and the Montreal Protocol on Substances That Deplete the Ozone Layer (see United Nations Environment Programme, 2000) could both be invoked to prevent extensive testing of stratospheric aerosol approaches. The UN Convention on Biodiversity (see Convention on Biological Diversity, 2011) has extremely broad language that some might interpret as preventing deployment of potentially harmful geoengineering systems. The Convention on the Prohibition of Military or Any Other Hostile Use of Environmental Modification Techniques (1976) might be interpreted as prohibiting geoengineering deployments that inflict severe harm to other nations.

It seems clear that these existing agreements, if vigorously enforced, could slow or halt initial attempts to test geoengineering systems. The ocean dumping and biodiversity agreements both have already been used for this purpose, as could the air-pollution conventions in the future. But it seems similarly clear that these existing agreements provide a thoroughly inadequate framework for governing geoengineering testing and deployment. None of these agreements explicitly addresses geoengineering, so there exists no consensus on what types of geoengineering activities they might allow or prohibit. A nation convinced that geoengineering provided potentially vital protection against the risks of climate change would be unlikely to accept these previous agreements as a *de facto* ban on research and testing.

Any cooperative deployment of geoengineering would require a level of agreement on issues far outside the scope of these existing agreements. Nations would have to allocate responsibilities for building, operating, and monitoring the system. They would have to agree on procedures for controlling the earth's thermostat—that is, determining the level of intervention, how it is adjusted over time, and procedures for shutting it down if something goes wrong (and even defining wrong). Nations would have to agree on appropriate compensation for those excessively harmed by the deployment. Victor (2008) attempts to capture this range of issues by contrasting any actual deployment with current speculations about potential systems. A real system, with adverse regional effects and leaving some impacts untouched, would, in practice, require a large portfolio of interconnected interventions, which Victor calls cocktail geoengineering. For instance, a sulfate aerosol intervention would not reduce ocean acidification. Thus, key ecosystems, such as coral reefs, might be protected by deacidifying regional waters (e.g., House et al., 2007). Regions that lose their rainfall might significantly upgrade their water-management systems, perhaps assisted by outside investment from those most benefiting from the aerosol intervention.

In this context, any near-term negotiations over geoengineering would, at best, likely prove inconclusive. Governments generally give environmental treaties extensive review and consideration. When ratified, such treaties usually represent a careful judgment of what commitments and obligations a country is sure it can meet (Chayes and Chayes, 1998). When treaty negotiations confront overwhelming uncertainty or when little agreement exists on how to frame the problem at hand, treaty negotiations often end in stalemate or produce only vague language that generates the illusion of agreement without any actual obligations. At present, in

⁴ Convention on the Prevention of Marine Pollution by Dumping of Wastes and Other Matter, 1972 and 1996 Protocol Thereto.

⁵ Convention on Long-Range Transboundary Air Pollution.

any negotiations seeking to establish international law governing geoengineering, some industrially advanced and technically sophisticated nations would seek to preserve their option to deploy such a system, believing that it might prove necessary in some future situation. Many nations, worried about the potential environmental consequences and believing that they were unable to deploy such a system themselves, might lean toward a ban. The details of Victor's cocktail geoengineering, the mix of compensation and corrective measures that might balance the needs of nations deploying geoengineering with those that could suffer some adverse consequences, could not contribute to a near-term deal because they remain so speculative for the foreseeable future. Given such obstacles, no comprehensive agreement governing geoengineering is likely any time soon. In addition, there seems to be little interest in adding such a controversial topic to the international climate-change agenda before an agreement on emission reductions is completed.

Influenced by such consideration, many commentators advocate a more sequential, bottom-up approach to geoengineering governance (Victor, 2008; Lin, 2009; Keith, 2010). In the near term, scientific societies and government research agencies would cooperate in establishing procedures, rules, and ethical standards governing research on and testing of geoengineering systems. The Royal Society, in support of the finding quoted at the start of this report, recommends that international scientific societies establish voluntary standards promoting transparency, appropriate validation and monitoring, public involvement, and de minimis standards for regulation of any geoengineering research. Cicerone (2006) recommends that scientific leaders propose a moratorium on large-scale field experiments (while supporting phase I research) until a system is in place to oversee the planning and implementation of such interventions. Blackstock and Long (2010) suggest that the scientific community determine "best practices" for research and restrict field testing until there is a broadly accepted international process for approving climatic impact tests. Keith, Parson, and Morgan (2010), particularly concerned that premature regulations (including bans on testing) should not unduly hamper research, propose a looser structure, in which informal consultations allow a governance structure to emerge slowly as research advances. Victor envisions a bottom-up process among scientists and government research agencies that will build a foundation of widely accepted, shared practice that can eventually be codified into formal norms for tests of geoengineering systems. He advocates decentralized geoengineering assessments led by different nations, emphasizing transparency and cooperation among them. These assessments will become the main source of useful information about geoengineering while helping the practitioners to develop and disseminate informal research norms for managing such tests.

Over time, the international network of scientists conducting geoengineering laboratory research and tests, the research norms and institutions they help develop to regulate such tests, and the knowledge generated could help provide the foundation for future international agreements governing the deployment of any geoengineering systems.

Such international networks of experts, which the international relations literature calls epistemic communities, often play a key role in policy coordination among governments (Haas, 1992) and the creation of binding international regulations (Slaughter, 2004). In the mid-1970s, biologists first agreed to defer certain experiments transferring genes between species and then developed ethical standards and guidelines governing such research. International networks of atmospheric scientists helped lay the foundation for the 1989 Montreal Protocol that banned substances harmful to the ozone layer. In the late 1960s, cooperation and discus-

sions among Soviet and U.S. nuclear scientists contributed to the 1972 Anti-Ballistic Missile Treaty.⁶

A process of bottom-up research-norm creation by scientific communities codifying practice that ultimately becomes institutionalized in international institutions or law might well prove the most-promising approach for governing research, testing, and, if necessary, deployment of geoengineering systems. The approach seems promising both on its own merits and in comparison with the clear weaknesses of alternative options. However, the strengths and weaknesses of such a research norm-creation policy have not been subject to any rigorous and systematic evaluation, including potential impacts on deployment decisions and likelihood of successful results. The remainder of this report provides an initial attempt at such an analysis.

⁶ Treaty Between the United States of America and the Union of Soviet Socialist Republics on the Limitation of Anti-Ballistic Missile Systems (1972).

A Vulnerability-and-Response-Option Analysis Framework for a Risk Assessment of Geoengineering Governance

Policies toward geoengineering, as with other potential responses to climate change, are best evaluated within an iterative risk-management framework (Bernstein, Pachauri, and Reisinger, 2008; NRC, 2010b). Typically, one can conduct risk management by estimating the probability and potential impacts of various events, evaluating the ability of alternative policy interventions to manage the risks by reducing either the probability or the impacts, and thereby comparing the efficacy of alternative risk-management strategies.

Geoengineering policy is not, however, amenable to this type of risk analysis. Similarly to many other climate-related decisions, geoengineering presents conditions of deep uncertainty, in which parties to a decision do not know or agree on the model linking actions to consequences nor the probability distributions for the key input parameters to these models (Lempert, Popper, and Bankes, 2003). Many authors have argued that the technical uncertainties surrounding geoengineering are currently so large as to defy any meaningful, quantitative risk analysis. But any judgments about appropriate policy toward geoengineering also depend on what are, at best, equally speculative expectations about the socioeconomic uncertainties surrounding the reactions to any research and deployment of such systems. Although enhanced physical models validated by experiments might reduce some technical uncertainties,¹ the risks associated with geoengineering will remain deeply uncertain because social, political, and economic factors are less amenable to prediction. In addition, judgments about appropriate geoengineering policy will depend on the wide range of values and ethical frameworks that different people will bring to these questions.

Under such conditions, geoengineering risk management might best be implemented within a vulnerability-and-response-option framework. Rather than estimate the impacts and likelihood of various events, such an approach would begin with one or more potential geoengineering policies, identify the future conditions that might cause such a policy to fail to meet its goals (as defined by one or more parties to the decision), and use the results of this vulnerability analysis to compare alternative policies or to suggest ways in which policies might be modified to reduce their vulnerabilities.² This approach might prove valuable for a wide range of climate-related decisions under conditions of deep uncertainty (NRC, 2010b; Lempert,

¹ Reduce but not eliminate. It is likely that large uncertainties will remain in spite of extensive research and development (R&D) until deployment and long-term observation of performance. Some researchers have argued that there is an irreducible uncertainty in climate prediction (Knutti, 2008; Dessai et al., 2009). Thus, the costs and potential benefits of R&D must be thoroughly examined.

² This report does not explicitly address the ethical issues involved with geoengineering. Nonetheless, any ethical inquiry on this topic should grapple with the potential consequences of alternative choices regarding such systems.

Nakicenovic, et al., 2004) but, with a few exceptions (e.g., Goes, Keller, and Tuana, 2010), has not been implemented for geoengineering.

Useful vulnerability-and-response-option analyses can be conducted using either qualitative or quantitative approaches. As an example of the former, the assumption-based planning (ABP) approach (Dewar, 2002) asks decisionmakers to identify key assumptions underlying an organization's plans. A key assumption is one whose failure would require changes in the plan and one that could plausibly fail during the lifetime of the plan. Lempert (2009) uses ABP to qualitatively compare the potential vulnerabilities of alternative international frameworks proposed for the Copenhagen negotiations in December 2009 and suggested a potentially more-robust alternative.

RDM provides a quantitative decision analytic framework for implementing a vulnerability-and-response-option analysis. In brief, RDM runs computer simulation models over thousands to millions of cases to project the performance of proposed policies over many plausible future states of the world, in which each state is described by different assumptions about model structures and different values of the input parameters to those models. RDM then uses statistical cluster analysis over the resulting database of model runs to characterize as simply as possible the conditions under which policies fail to meet their performance goals (Bryant and Lempert, 2010). Analysts and decisionmakers can then use the resulting scenarios to identify and evaluate potential responses to these vulnerabilities. This RDM approach has been used to address a wide variety of policy challenges, such as water management under climate change, responses to abrupt climate change, terrorism insurance, flood control in New Orleans, and energy policy, and has been used both as analytic tool and as a means to facilitate discussions among policymakers ranging from local water managers to cabinet members in a national government.³

This study conducted the initial stages of an RDM analysis on the challenge of governing geoengineering research. The current work is to be considered from the prospective of a policymaker whose only concern is the ability to utilize geoengineering in a crisis, not the broader question of the fundamental value of additional knowledge of how the climate behaves.

The analysis constructs a simple simulation model that relates choices by the U.S. government regarding alternative governance approaches for geoengineering to various consequences for the United States and the world, contingent on the outcomes of several key uncertainties. The analysis then uses this simulation to compare the performance of each alternative policy over two dozen cases, each case represented by a different combination of assumptions about the performance of geoengineering technologies and the sociopolitical environment in which they might (or might not) be deployed, and uses these results to characterize some of the vulnerabilities—defined as scenarios in which a proposed policy performs significantly worse than other options. We then offer a qualitative discussion of some of the implications of these potential vulnerabilities.

This simple analysis rests between a qualitative ABP and a fully quantitative RDM analysis. It proves useful for several reasons. First, it demonstrates the potential ability for a risk analysis based on a vulnerability-and-response-option analysis to inform the debate on geoengineering. Second, it helps define the steps needed to conduct a full RDM analysis to address such governance issues. Third, it provides some intriguing, if only suggestive, policy results.

³ The RAND Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition (2010) website provides a summary of many RDM applications.

Key Factors Considered in the Analysis

This report pursues its assessment of the risks of various approaches to geoengineering governance with a state-centric analysis focusing on the interplay of two actors—the United States and an unspecified other country—either of which could unilaterally implement geoengineering if it judged that doing so would be in its best interest. This focus addresses the strong correlation among the countries that have the financial and technical capabilities to deploy geoengineering and those that emit large quantities of greenhouse gases. An important risk of geoengineering is that countries might use the potential for such systems as a rationale to avoid taking actions to reduce greenhouse gases. We assume that the United States and others have similar technical capabilities (there is no difference in the success rate of tests or deployments) and comparable international behavior, with the only difference being that the United States is more likely to abide by research norms and international agreements than others. This latter assumption rests on the notion that, as a relatively open society, the United States often regards itself, rightly or wrongly, as more likely than some other nations to abide by its legal obligations. A more-extensive analysis might explore more thoroughly the extent of the asymmetry among nations in abiding by their treaty obligations.

This state-centered focus neglects, however, other potentially important actors that might affect the development and deployment of geoengineering. For instance, the analysis does not explicitly consider the role of scientific (epistemic) communities in setting and enforcing norms regarding geoengineering. Nor does the analysis consider the role of other nonstate actors, such as for-profit firms, that might take steps that would make deployment of geoengineering systems more or less likely (Blackstock and Long, 2010). The potential impacts of scientific communities and other nonstate actors enter the analysis in the most-general way only when considering cases in which governments' actions do not turn out as intended (e.g., a government might favor geoengineering tests but such tests do not occur). This analysis thus does not examine some important policy questions currently under debate, such as whether governments or scientific communities should take the lead in developing norms. Rather, the analysis addresses the question of what risks are associated with a U.S. government policy of encouraging the development of norms, irrespective of the forum. Clearly, some risks might depend on the choice of forum, but we believe that there is utility in addressing the more-general question, both because it provides some useful policy insights and because it demonstrates a risk-assessment methodology that could be more-broadly applied.

As is often the case with RDM analyses, it is useful to group its elements into four categories: policy levers, uncertainties, measures of merit, and relationships (Lempert, Popper, and Bankes, 2003). Policy levers (L) are near-term actions that, in various combinations, comprise the strategies decisionmakers wish to explore. This choice helps frame the analysis. In particular, it requires identifying the policymakers who constitute the main audience. The choice of measures and the difference between what factors constitute levers and which constitute uncertainties depends strongly on who is presumed to be making the decisions. Exogenous uncertainties (X) are factors, outside the control of the decisionmakers, which could nonetheless prove important in determining the success of their strategies. Measures (M) are the performance standards that decisionmakers and other interested communities would use to rank the desirability of various outcomes. Relationships (R) describe, using the measures, the outcomes that would result from choosing a particular set of policy levers contingent on the uncertainties.

Table 3.1 lists the levers, uncertainties, and measures used in this analysis. We now describe each of the four factors (X, L, M, and R) in turn.

Levers

The presumed audience for this analysis is the U.S. government, which, over the next few years, will make decisions about the type of international governing framework for geoengineering it believes should be put in place. We consider three alternative options (as shown in the upper right-hand corner of Table 3.1): Strong Norms, Ban, and No Norms, which we describe as follows:

- **Strong Norms:** The U.S. government encourages other governments and scientific societies to promote research norms for the conduct of geoengineering research. Such norms describe standards for protecting against adverse environmental impacts from regional or global tests, suggest compliance with existing applicable treaties, and might provide guidance on provisions for compensation for any adverse impacts of tests. Having strong research norms also encourages open sharing of all experimental data and results. As one important feature, we assume that having strong research norms promotes a noninterference standard that precludes conducting more than one large-scale test at any one time. Without such a standard, it could prove almost impossible to disentangle the impacts of multiple, simultaneous global tests of geoengineering.⁴
- **Ban:** The U.S. government encourages other governments and scientific societies to promote bans that prevent any type of geoengineering research.
- **No Norms:** The U.S. government seeks to prevent the emergence of any type of norms or international agreements governing the testing of geoengineering.

This set of stylized policy options helps demonstrate the methodology and offers some suggestive policy results. However, the three present only a limited representation of the full slate of alternatives facing U.S. policymakers.

This report's state-centric focus imposes important limits on the policy questions it can address. First and foremost, the report does not explore the appropriate forum for developing norms, whether they are best developed by governments in a formal negotiating process (see

Table 3.1
Factors Considered in the Analysis: Exogenous Uncertainties (X), Policy Levers (L), Measures (M), and Relationships (R)

X Severity of climate change (catastrophic, severe, or mild) Technical potential for geoengineering (likely or unlikely) Potential for agreements (favorable or unfavorable) Influence of geoengineering on emissions (strong or weak)	L Strong Norms: The United States promotes strong norms. Ban: The United States works to institute a ban. No Norms: The United States works to prevent the establishment of any norms.
R Bayesian network	M Impacts on the United States Impacts on others Potential for conflict Collateral damage (e.g., ecosystems)

⁴ This assumption has a significant impact on our assessment of the vulnerabilities of alternative policies.

Barrett, 2008), through existing international bodies, or more informally by scientific societies (see Victor, 2008). Rather, the analysis assumes that a decision by the U.S. government to favor or oppose the setting of such norms will make their development, in whatever forum, more or less likely. Similarly, the analysis is not well configured to examine how the U.S. government might pursue hybrids among the three options considered here. For instance, the United States could pursue different approaches for large- and small-scale geoengineering testing. The United States could also follow a wait-and-see strategy, which allows others to become first actors on this issue. A wait-and-see strategy (balanced against the risk of delaying action) might permit the United States to better access both the certainty of future climate predictions and the risks of geoengineering strategies before deciding on a policy option. In this report, the United States can only take a strong stand one way or the other.

It is also important to note that this report focuses on near-term decisions about the rules governing geoengineering development. As described later, the analysis does consider cases in which the international community makes and does not make international agreements on the deployment of geoengineering systems. However, the analysis regards such agreements as future events outside the direct control of today's decisionmakers. Today's decisionmakers can influence such events only via the preferences they express today regarding the international framework for governing geoengineering. This is consistent with our belief that decisions on research will affect future deployment decisions.

Uncertainties

Once the U.S. government has chosen among the Strong Norms, Ban, and No Norms policies, many uncertainties will come to affect the potential success or failure of this choice. This analysis focuses on four as particularly germane (upper left-hand corner of Table 3.1): the potential for serious adverse climate change, the potential for successful geoengineering technology, the potential for international agreement on actions to address climate change, and the extent to which progress on geoengineering will influence (undermine) progress on reducing greenhouse gas emissions. These uncertainties form the basis of the small set of scenarios we examine in the current analysis. Here, we provide a qualitative description of the range of uncertainties. Appendix C describes how we quantify these ranges.

Potential of Geoengineering Technology. The potential for geoengineering to counteract the effects of climate change is deeply uncertain. We summarize these uncertainties with two states:

- **Likely to succeed:** In this state, a competent nation is likely to be able to develop and deploy a geoengineering system that significantly reduces the adverse impacts of climate change either regionally or globally.
- **Unlikely to succeed:** In this state, any deployed geoengineering system is likely to produce consequences more adverse than the climate change it was intended to counter, either regionally or globally.

Potential Impacts of Climate Change. The potential for serious adverse impacts from climate change depends on the future level of greenhouse gas emissions, as well as on uncertain properties of the climate system, such as the climate sensitivity, which measures how much the climate changes due to a given level of human emissions, the presence of any instabilities (tip-

ping points) in the climate system, and the response of weather patterns, ecosystems, and other factors to the changes in climate. We summarize these uncertainties with three states:

- **Mild:** In this state, the climate is not very sensitive to human emissions, and human society and ecosystems can easily adapt to any consequences from climate change.
- **Severe:** In this state, the climate is highly sensitive to human emissions, and any climate changes will seriously damage ecosystems and cause social and economic damage beyond that to which human society can easily adapt.
- **Catastrophic:** In this state, the climate is highly susceptible to extreme tipping points that cause sweeping disruptions to human society and devastate the world's ecosystems.

Potential for Agreements. Any efforts at reducing greenhouse gas emissions, as well as deploying geoengineering systems, will prove more effective if the world's nations reach agreement on how best to proceed and how to allocate the responsibility and costs for implementing their plans. A wide variety of factors, ranging from the skill of national leaders to how fast the global economy is growing, can influence nations' ability to reach such agreements. We summarize these uncertainties with two states:

- **Favorable international conditions:** In this state, many factors combine to make nations eager and able to reach agreement on greenhouse gas reductions or the rules governing any deployment of geoengineering systems. Note that a favorable climate for cooperation does not necessarily lead to agreements on either geoengineering deployments or greenhouse gas reductions but makes it more likely that efforts to pursue such agreements will prove successful.
- **Unfavorable international conditions:** In this state, many factors combine to make it difficult for nations to reach such agreements.

Influence of Geoengineering on Efforts to Reduce Emissions. Serious efforts to test geoengineering systems could make it more difficult to reach and implement agreements on deep reductions of greenhouse gas emissions. We summarize this uncertainty with two states:

- **Strong influence:** In this state, efforts to reduce greenhouse gas emissions are significantly undermined by the intention to test geoengineering systems.
- **Weak influence:** In this state, efforts to reduce greenhouse gas emissions are not undermined by the intention to test geoengineering systems.

Measures

What considerations might the U.S. government use to assess the success or failure of its approach to geoengineering? We offer four measures (shown in the lower right-hand corner of Table 3.1) to summarize the criteria U.S. policymakers might use. These measures focus on the consequences of the deployment (or not) of geoengineering systems on climate change and other international security concerns. These measures include the following:

- **Impacts on the United States:** This measure reflects U.S. government concern with direct adverse climate impacts on the United States, such as drought, increased hurricanes, extreme weather events, or the spread of tropical diseases into North America. The net

impact on the climate is a consequence of the success or failure of geoengineering, the severity of the climate scenario, and the extent of greenhouse gas emission reductions. Greenhouse gas reductions, if strong, are judged to more-successfully reduce the adverse impacts of climate change than would a successful geoengineering deployment. The extent of success is moderated by the climate scenario—all other factors being equal.

- **Impacts on others:** This measure reflects U.S. government concern with direct adverse climate impacts on other countries sufficiently large and sophisticated to deploy their own geoengineering systems. The net climate impact is a consequence of the success or failure of geoengineering, the severity of the climate scenario, and the extent of greenhouse gas emission reductions. Greenhouse gas reductions, if strong, are judged to more-successfully reduce the adverse impacts of climate change than would a successful geoengineering deployment. The extent of success is moderated by the climate scenario—all other factors being equal.
- **Collateral damage:** This measure reflects U.S. government concern with global impacts from geoengineering, such as ozone depletion, destruction of natural ecosystems, failure to treat ocean acidification, and the direct impacts of climate change on states unable to deploy geoengineering systems. Collateral impacts occur if geoengineering fails or has significant adverse impacts on countries other than those that deploy it. We assume that, if geoengineering is implemented on less than a global scale, there will be adverse secondary impacts or untreated impacts of global warming. Weak efforts to reduce greenhouse gas emissions increase collateral impacts.
- **Conflict:** The United States will be concerned about the potential for international conflict over the deployment of geoengineering. We assume that a high potential for conflict results when a deployed geoengineering system fails—thereby causing significant damage to many nations, including many that might have had no direct role in the deployment—or when a geoengineering system generates asymmetric results—successfully reducing climate impacts for the nations that deploy it but failing to protect or potentially damaging other nations.

One might imagine a shorter or longer list of the best measures of the ultimate success of any near-term geoengineering policy. In particular, these measures and the underlying analysis do not address any consequences of developing geoengineering systems (such as field testing) other than the effect such efforts might have on the likelihood of reaching an international agreement on reducing greenhouse gas emissions and on the likelihood of deploying a geoengineering system.⁵ The analysis thus ignores any adverse consequences or conflicts engendered solely by the conduct of geoengineering research. Nonetheless, we believe that these four measures represent a sufficiently diverse set both to demonstrate the risk-assessment methodology and to provide a useful summary of important strengths and weaknesses of the alternative policies considered here. It is also useful to note that analysis only identifies the trade-offs among these measures and, as is typical with many multicriterion decision analyses, explicitly avoids any attempt to weigh their relative importance.

⁵ An expanded network, especially if time dependencies are considered, might include adverse impacts (either climatic or political) caused by the tests, independent of any deployment decision.

Relationships

This analysis uses a simple simulation model to link the measures to the levers and uncertainties—that is, to estimate the outcomes as described by each of the four measures that might result, contingent on the resolution of each of the uncertainties, from a U.S. decision to pursue each of the three alternative near-term policies. The particular type of model used here is called a Bayesian belief network (BBN).⁶ A BBN is a well-known mathematical construct that allows one to examine the complex interactions of a chain of uncertain events. It illuminates how the choice of near-term decision (for example, banning or not banning geoengineering research) might change the likelihood of various climate outcomes. Such a network consists of a collection of nodes, each having several possible states, and the links among the nodes. The state of any one node depends probabilistically on the states of the previous nodes to which it is linked.


A BBN simulation is useful for this study because it provides a simple framework for characterizing and quantifying the relationships among a wide variety of factors. Such Bayesian networks have similarly proved useful for addressing numerous types of decision problems, including many involved with environmental decisions and with political strategy (McCann, Marcot, and Ellis, 2006; Jackman, 2004; Marcot et al., 2001; Cain, 2001). However, using such networks does impose some limitations on the analysis. For instance, representing the detailed timing of events adds considerable complexity to a BBN, which would complicate any attempt to consider U.S. government policies that are explicitly designed to evolve over time in response to new information. A BBN has difficulty representing systems with multiple actors, such as international systems influenced by scientific communities, for-profit firms, or even self-organizing coalitions of states. Addressing such factors as the timing of policies or multiple actors might require a different type of simulation model to express the relationships among the uncertainties, levers, and measures. It is important, however, to note the important distinction between a simulation model and the decision analytic framework used to exploit the information in that model. The vulnerability-and-response-option framework presented in this report offers a decision analytic approach that can be used with a wide variety of different types of models.

As shown in Figures 3.1, 3.2, and 3.3 and described in more detail in Appendixes B and C, the network used in this analysis consists of 17 nodes: one for the policy levers, four for the uncertainties, four for the measures, and eight describing the relationships among them. These relationship nodes are usefully clustered: One node relates to any reduction of greenhouse gas emissions, three represent the process of any field testing of geoengineering systems, and four relate to the process of deploying any geoengineering system.

For instance, the node representing the United States' decision to deploy a geoengineering system (labeled "USDeploys" in Figures 3.1–3.3) has states representing a multilateral decision to deploy, a unilateral decision to deploy, or a decision not to deploy. The probability of the node adopting one of these three states depends on the states of previous nodes representing the severity of climate change, whether the international community has begun serious emission reductions, any international agreements on deployment, and the results of any tests of geoengineering systems. The node for emission reductions (labeled "EmissionsReductions" in Figures 3.1–3.3) has states representing strong and weak reductions. The probability that the

⁶ We implemented this simulation using Netica software from Norsys Software Corporation.


Figure 3.1
Simulation Nodes for Strong Norms Policy



NOTE: This schematic of a simulation shows Strong Norms policy in a scenario with geoengineering potential likely, favorable international conditions for emission reductions, potentially catastrophic impacts from climate change, and high influence of geoengineering tests on emission-reduction efforts. GE = geoengineering.

RAND TR846-3.1


Figure 3.2
Simulation Nodes for Ban



NOTE: This schematic of a simulation shows a Ban policy in the same scenario shown in Figure 3.1.

RAND TR846-3.2

Figure 3.3
Simulation Nodes for No Norms Policy



NOTE: This schematic of a simulation shows a No Norms policy in the same scenario used in Figures 3.1 and 3.2.

RAND TR846-3.3

U.S. deployment node is in the state “NoDeployment” might be higher if the EmissionsReductions node is in the state “Strong.” The full network follows such dependencies along the full range of steps linking the initial policy decision (levers) to the ultimate outcomes (measures) contingent on the state of the uncertainties.

As illustrated in Figures 3.1, 3.2, and 3.3, the U.S. government’s policy choice regarding research norms affects subsequent decisions about testing and the international community’s ability to reach any international agreement governing geoengineering deployments. International conditions affect potential agreements about emission reductions and on rules governing any geoengineering deployments. Decisions about setting norms and testing affect subsequent decisions regarding emission reductions. In addition, the influence of geoengineering on emissions affects nations’ ability to come to agreement and implement significant reduction in greenhouse gases. The state of the climate, the results of testing, and the status of emission reductions all influence deployment decisions. The potential for geoengineering affects the success of any tests and deployment of geoengineering systems. The success of any geoengineering deployments and the aggressiveness of any emission reductions affect the climate impacts experienced by the United States and other nations. Any potential conflict is affected by the disparity in such impacts.⁷

This simulation is specifically designed to compare the implications of the three alternative near-term policy options—Strong Norms, Ban, and No Norms—over a wide range of scenarios, each constructed from different combinations of the four types of uncertainty. As one example, this simulation can reproduce each of the three cases in which nations might deploy a geoengineering system described in “Geoengineering Deployments and Their Risks” in Chapter Two.

The model represents the case in which one or more nations deploys in response to others’ failure to reduce emissions as one in which unfavorable international conditions for agreements (center node on the left of the three figures) leads to failed efforts at emission reduction (upper node in the middle of the three figures), and severe climate change (upper node in the middle of the three figures) leads to U.S. deployments. The model represents the case in which one or more nations deploys in response to catastrophic climate change as one in which such a change leads to geoengineering deployments despite a successful deployment agreement. The model represents the case in which one or more nations deploys to avoid making significant emission reductions as one in which successful geoengineering tests (cluster of three nodes in the lower left of the three figures), combined with a strong influence of geoengineering on emission reductions (node in the upper left-hand corner of the three figures) leads to weak emission-reduction efforts and to geoengineering deployments.

Note that the structure of the network embodies certain assumptions about the timing of events. We assume that tests of geoengineering systems begin and any agreements about the rules for deploying such systems occur prior to any implementation of strong policies to reduce greenhouse gas emissions. Thus, such testing and agreements can influence choices about such reduction policies. We assume that decisions about emission-reduction policies must be made before much more is learned about the state of the climate system. We also assume that any decision about whether to deploy a geoengineering system will occur after any emission-reduction policies have been implemented (or not) and more is learned about the potential severity of

⁷ Adverse climate impacts of field tests have not been considered in this simplified model.

climate change. These assumptions are consistent with the notion that SRM geoengineering technologies represent a contingency option (to be used either multinationally or unilaterally) that would be seriously considered only if the impacts of climate change were likely to be severe and there were little progress on emission reductions, or to avoid or recover from abrupt climate change.

Appendixes B and C provide more-detailed descriptions of each of the relationships (R) and their potential states, and the probability tables that link them. The network's results are driven by the choice of links and the choice of values to fill the probability tables. Although the causal links are straightforward, this analysis employs pure subjective probability estimates, assigned by the authors using our best judgments and with a desire to illustrate the potential interactions and interdependencies of the levers, uncertainties, and measures of merit.

To increase consistency and capture correlations among the probability judgments, we expressed each probability as a function of several relevant factors for each node (called influence factors). For instance, we assumed that the probability that the international community would agree to rules governing the deployment of any geoengineering system can be reduced depending on the international conditions and the near-term decision on research norms, and then we separately estimated the impact of each on the probability of a deployment agreement. We then multiplied these factors together to give the probabilities for the six different cases (three policies times two conditions). The details for this and the resultant complete probability tables are shown in Appendix C. Clearly different probability estimates could yield different policy vulnerabilities. We use the network as an exploratory tool within a scenario-based approach specifically designed to be exceedingly mindful of the limitations of our simulation.

Caveats

Before presenting results in the next chapter, it is important to reemphasize what this analysis will not do. First, the simulation is, in no sense, predictive. Rather, it provides a what-if mechanism that allows us to systematically follow the implications of alternative chains of assumptions. It is well known that even experts have trouble following in their heads a string of logical relationships over more than a few causal steps. The simulation provides a formal accounting system that allows us to overcome these barriers and to systematically compare the different and sometimes unexpected consequences of various combinations of assumptions. Bankes (1993) distinguishes between consolidative and exploratory models. The former can be validated and thus used in predictive analysis. The latter can still help decisionmakers evaluate the strengths and weaknesses of alternative responses to deeply uncertain futures. The simulation here is very much an exploratory model.

Second, this analysis in no way represents the full decision analytic treatment possible with such an exploratory model. A full RDM treatment would vary over all the assumptions in the simulation and perhaps consider different structural or expanded forms for the Bayesian network. A full analysis would run many thousands of cases and conduct statistical cluster analyses to characterize the vulnerabilities of the three policies (Bryant and Lempert, 2010). In particular, as described above and in Appendixes B and C, the analysis' simulation model uses probabilities as its input parameters, each representing the likelihood that one state—for instance, a successful test of a geoengineering system—will lead to another state—for instance, the deployment of such a system. Here, we consider only a small set of discrete values for each

probability, corresponding to the cases shown in Table 3.2.⁸ A full RDM analysis would regard these probabilities as imprecise—that is, assuming that any of a multiple set of plausible values is possible (Lempert and Collins, 2007).

Having systematically identified scenarios in which the Strong Norms, Ban, and No Norms policies might fail to achieve the U.S. government’s goals, the full RDM analysis would then use this information to design new policies that might ameliorate some of these vulnerabilities. This step would generate trade-off curves that could help decisionmakers choose the policy they felt was most robust. Typically, such trade-off curves suggest probability thresholds for key scenarios. Such thresholds represent the probability that a scenario would need to have in order to justify abandoning a proposed policy in favor of some other option (Groves and Lempert, 2007; Lempert, Groves, et al., 2006). The analysis could then explore how the alternative policies might be modified to make them more robust—that is, less sensitive to the identified vulnerabilities. The analysis here provides only a much-simplified first cut at such an analysis, focusing on only a few key uncertainties and providing only a qualitative description of potential improvements to the initial menu of policy options and the trade-offs among them. This simplified analysis will identify key scenarios but explore such probability thresholds only qualitatively.

Evaluating Policies Across Multiple Scenarios

We can now use this simulation model to evaluate how the measures in Table 3.1 depend on the U.S. government’s choice of policy and how the uncertainties are resolved. Such an analysis needs to address what combinations of uncertainties it will consider—that is, specify the experimental design used to explore the input parameters to the model. In addition, the analysis needs to provide a way to succinctly compare the performance of alternative policies across many scenarios. This discussion will address both these issues.

This study employed a relatively simple experimental design. As shown in Table 3.1, the analysis considers four uncertain socioeconomic and physical factors: the potential for international agreements, the influence of geoengineering systems on emission reductions, the seriousness of future climate change, and the potential of geoengineering technology. As summarized in Table 3.2, the analysis considers three possible values for the seriousness of future climate change and two possible values for the other three uncertainties. To construct the scenarios, we

Table 3.2
Range of Cases Considered for Each Uncertain Parameter

Socioeconomic Systems		Physical Systems	
Potential for Agreements	Influence of Geoengineering on Emission Reductions	Seriousness of Climate Change	Geoengineering Technology Potential
Favorable Unfavorable	Strong Weak	Catastrophic Severe Mild	Likely Unlikely

⁸ This study made this simplification primarily because it lacked sufficient resources to build the software harness needed to conveniently run many cases of the model with different assumptions about the input probabilities.

consider all possible combinations of the four uncertain parameters, which yields 24 (3×23) cases.

To help summarize the performance of the three policies, we calculate how far in each scenario for each measure of merit the performance of a given strategy deviates from the performance of the optimal strategy in that scenario. We calculate this deviation using an expression called the relative regret, given by

$$\text{Rel_Regret}_{l,x,m} = \frac{\text{Regret}_{l,x,m}}{\text{Min}_{l' \in L}(M_{l',x,m})}, \quad (1)$$

where

$$\text{Regret}_{l,x,m} = M_{l,x,m} - \text{Min}_{l' \in L}(M_{l',x,m})$$

is the regret introduced by Savage (1954) to implement the mini-max decision criterion and where $M_{l,x,m}$ is the performance of policy l , in state of the world x , according to measure m . We use a regret measure for this analysis in order to focus on the comparative performance of alternative policies and thus the vulnerability of policies over a very wide range of plausible future conditions. In some scenarios—for instance, one with catastrophic climate change and unfavorable conditions for any international agreement—the future will prove difficult no matter what policy choices are made today. In other scenarios—for instance, one with mild climate change—the future will prove benign unless policy choices go significantly awry.⁹ For any given policy, the regret will be small when the policy performs close to the best one can possibly perform in that scenario, whether the scenario is benign or dire. The regret is large when the policy performs far worse than the best possible in that scenario. We use relative regret, as opposed to simple regret, because we assume that a given nonzero deviation from optimality will matter more in a benign scenario, in which the value of

$$\text{Min}_{l' \in L}(M_{l',x,m})$$

is low, than it will in a dire scenario, in which the value of

$$\text{Min}_{l' \in L}(M_{l',x,m})$$

is very high.

As an example of how this regret measure summarizes the performance of alternative policies, consider the scenario represented by the model configuration shown in Figures 3.1, 3.2, and 3.3, in which, as shown in the boxes representing uncertainties (labeled with X's), geo-engineering technology is likely to succeed, climate change can be catastrophic, geoengineer-

⁹ Note that the simulation includes no cases in which the pursuit of geoengineering or other policy choices can cause adverse outcomes in a mild-climate scenario. If geoengineering could truly be unilaterally deployed, it is conceivable that a nation might pursue such systems to correct regional environmental issues or, for example, enhance regional agricultural production, even with a mild global climate impact. An expanded RDM analysis could examine these scenarios.

ing activities can significantly affect the commitment to emission reductions, and international conditions are favorable for agreements. Table 3.3 shows the performance for each of the three policies for each of the four measures in this scenario. These values are drawn, as discussed, from the boxes in Figures 3.1, 3.2, and 3.3, representing measures (labeled with M's). Note that a lower score indicates better performance, so the Strong Norms policy performs best for three of the measures—impacts on the United States, impacts on others, and collateral damage—while the No Norms policy performs best for the conflict measure.

Table 3.4 shows the relative regret for each policy for each measure calculated using Equation 1. Strong Norms remains best (0 percent regret) for the first three measures. No Norms remains best for the conflict measure, but, with 14 percent regret, Strong Norms performs nearly as well. Overall, the regret values provide a convenient measure of how well each policy performs compared to the best possible in this scenario for each measure.

Figure 3.4 summarizes results, such as those in Table 3.4, for 16 of the 24 scenarios—those with severe and catastrophic climate change. Each row of the table represents a single scenario. The first four cells in each row show the relative regret of choosing the Strong Norms policy for each of the four measures of merit. The next four cells show the relative regret of choosing the Ban policy, and the last four cells show the relative regret of choosing the No Norms policy. We have shaded the results to make it easier to see patterns across the policies and scenarios. Green indicates regret less than 10 percent, red indicates a regret greater than 25 percent, and yellow indicates a regret between 10 and 25 percent. These thresholds are chosen to aid in presentation only and have no effect on the results. Furthermore, although the results are shown to two significant figures, the numbers are given for comparative purpose only and do not indicate significance to that level. Note that Figure 3.4 shows only those scenarios with severe or catastrophic climate change because, if climate change is mild, the choice of geoengineering policy proves unimportant (at least at this level of analysis). In all mild-

Table 3.3
Performance for Three Policies for the Scenario Shown in Figure 3.1

	Impacts on the United States	Impacts on Others	Collateral Damage	Conflict
Strong Norms	46	49	52	48.4
Ban	91	57	100	76.3
No Norms	80	80	90	42.6

NOTE: Scores, dimensionless numbers ranging from 0 to 100 with lower values more desirable, are taken from values in the Measures nodes in Figures 3.1–3.3.

Table 3.4
Relative Regret for Three Policies for the Scenario Shown in Figure 3.1

	Impacts on the United States	Impacts on Others	Collateral Damage	Conflict
Strong Norms	0	0	0	14
Ban	99	16	92	79
No Norms	76	64	74	0

NOTE: Scores are relative regret (%) calculated from the values in Table 3.3 using Equation 1.

Figure 3.4
Relative Regret (%) of Three Policies Over 16 Scenarios

GE Potential	Climate	GE Influence	Potential for Agreements	Strong Norms				Ban				No Norms			
				Impacts to U.S.	Impacts to Others	Collateral Damage	Conflict	Impacts to U.S.	Impacts to Others	Collateral Damage	Conflict	Impacts to U.S.	Impacts to Others	Collateral Damage	Conflict
Likely	Catastrophic	Strong	Favorable	0	0	0	14	99	16	92	79	76	64	74	0
Likely	Catastrophic	Strong	Unfavorable	0	22	0	67	26	0	13	65	17	37	9	0
Likely	Catastrophic	Weak	Favorable	0	0	0	0	227	77	241	158	148	100	125	47
Likely	Catastrophic	Weak	Unfavorable	0	26	1	53	35	0	19	62	6	27	0	0
Likely	Severe	Strong	Favorable	0	9	0	8	75	0	68	0	70	80	89	13
Likely	Severe	Strong	Unfavorable	0	19	0	33	62	0	35	37	22	37	13	0
Likely	Severe	Weak	Favorable	0	9	0	0	68	0	38	39	40	38	49	71
Likely	Severe	Weak	Unfavorable	0	16	0	39	59	0	28	41	9	20	0	0
Unlikely	Catastrophic	Strong	Favorable	0	8	6	148	5	0	0	95	3	6	0	0
Unlikely	Catastrophic	Strong	Unfavorable	0	6	4	165	3	0	0	92	3	4	0	0
Unlikely	Catastrophic	Weak	Favorable	54	67	72	149	52	39	47	92	0	0	0	0
Unlikely	Catastrophic	Weak	Unfavorable	19	25	27	164	17	10	16	90	0	0	0	0
Unlikely	Severe	Strong	Favorable	36	44	44	93	0	0	0	0	34	44	39	35
Unlikely	Severe	Strong	Unfavorable	12	24	21	95	0	0	0	21	0	8	3	0
Unlikely	Severe	Weak	Favorable	15	21	22	69	0	0	0	0	8	16	14	42
Unlikely	Severe	Weak	Unfavorable	23	26	28	105	12	4	8	29	0	0	0	0

NOTE: Green indicates relative regret less than 10 percent, red greater than 25 percent, and yellow between 10 and 25 percent.

RAND TR846-3.4

climate scenarios, all three alternative options perform identically for all four measures because we assume that no nation would deploy a geoengineering system unless it were experiencing serious impacts from climate change. With a mild climate, no nation deploys a geoengineering system and no nation suffers any impacts from climate change irrespective of the near-term U.S. choice of geoengineering policy.

To better understand the information in Figure 3.4, consider its first row, which summarizes the scenario considered in Table 3.4. As shown in that table, the Strong Norms policy in this scenario has 0 percent regret for three of the four measures: impacts on the United States, impacts on others, and collateral damage. The policy's regret for the conflict measure is, however, 14 percent. That is, for the first three measures, the Strong Norms policy performs best in this scenario. For the conflict measure, the policy generates a score for conflict more than 14 percent larger than does the No Norms policy in this scenario.

Comparing the network diagrams in Figures 3.1, 3.2, and 3.3 helps to explain these results. Under the Strong Norms policy, the United States leads a multilateral testing program for geoengineering systems, as seen in Figure 3.1 by the 100-percent probability of the *USTests* node being in the *Multilateral* state. Because, in this scenario, geoengineering technology is likely to succeed, the tests prove favorable, as seen by the 0-percent probability of the *TestResults* node being in the *Failure* or *NoResults* state. When climate change proves catastrophic in this scenario, nations deploy geoengineering systems with a high probability of success, as seen in the *DeploymentResults* node. Unfortunately, this scenario assumes that this geoengineering activity has a significant negative impact on emission-reduction activities, as seen by the high probability of the *EmissionsReductions* node being in the state *Weak*.

In comparison, the Ban policy prevents testing by the United States, as seen in Figure 3.2 by the 100-percent probability of the *USTests* node being in the *NoTest* state. But other nations might test, as seen by the approximately 75-percent probability of the *OthersTest* node being in the *Unilateral* state. Nonetheless, the Ban policy significantly reduces testing compared to the Strong Norms case so that emission-reduction efforts are much more likely to prove successful. When, nevertheless, climate change proves catastrophic, nations deploy geoengineering in an attempt to protect themselves. But these systems have a high probability of failure because we assume that, without tests, geoengineering is less likely to succeed. Compared to the Strong Norms policy, the Ban policy's high geoengineering failure rate weighs more heavily on the measures of merit than the Ban policy's larger potential for emission reductions. The latter policy performs worse in this scenario for all four measures, as seen by each of the four nodes at the far right of the figure, which show higher readings (and thus worse outcomes) in Figure 3.2 than in Figure 3.1.

The No Norms policy in Figure 3.3 performs worse in this scenario than the Strong Norms policy did for three of the four measures but performs better for the conflict measure (thus the yellow in Table 3.4). Under the former policy, both the United States and other nations are likely to pursue geoengineering research independently. The conflicting test programs are more likely to fail. Nations are also unlikely to reach any agreement governing geoengineering deployment. When climate change proves catastrophic, nations are less likely to respond with what otherwise might have been successful geoengineering deployments.¹⁰ Thus, impacts on the United States, impacts on others, and collateral damage are worse with the No

¹⁰ Nations are more likely to deploy an untested technology than one that failed tests.

Norms policy than with the Strong Norms policy. But, without geoengineering deployments, the potential for conflict is reduced because no nation suffering adverse climate impacts can blame another's geoengineering system for its troubles.

Identifying the Vulnerabilities of Alternative U.S. Geoengineering Governance Policies

The simulation results shown in Table 3.4 and the first row of Figure 3.4 compare the performance of the three alternative U.S. government policies—Strong Norms, Ban, and No Norms—in a single scenario reflecting one particular set of assumptions about the future potential of geoengineering technology, the severity of climate change, the influence of geoengineering development on emission reductions, the potential for international agreements, and the relationship among these factors and U.S. policy choices. Because of the deep uncertainty surrounding all of these issues, the comparison among the policies in this single scenario is not in and of itself very useful. However, the essence of the vulnerability-and-response-option approach described in Chapter Three is to compare the performance of these policies over a wide range of many carefully chosen scenarios to understand the patterns in their vulnerabilities.

This chapter describes the results of such a comparison across scenarios.

The basic pattern seen in Figure 3.4 is straightforward. Strong Norms performs generally well in those scenarios in which SRM geoengineering technology is likely to succeed (first eight rows in Figure 3.4). Ban and No Norms perform generally well when SRM technology is unlikely to succeed (next eight rows).

Examining the simulation results suggest the reasons underlying this result. Strong Norms works well in the former cases because it increases the chances that, if adverse climate conditions demand a geoengineering deployment, a well-tested system is available. In the latter cases, Strong Norms performs less well than the alternative policies because it can slow efforts at emission reductions and can result in overconfidence—the deployment of a geoengineering system that is expected to work but does not.

Figure 3.4 also shows a basic pattern in the vulnerabilities of a U.S. policy promoting a Ban, which tends to work better than the other policies for other nations that feel free to break it (see the “Impacts to Others” column under “Ban” in the figure). This policy generally has high regret for three of the four measures when geoengineering is likely to succeed. But the policy has low regret under these conditions for the Impacts to Others measure because other nations might test in defiance of the prohibition and subsequently deploy viable systems that reduce climate impacts on themselves. The Ban policy performs poorly for other nations only when geoengineering is likely to succeed, climate change proves catastrophic, and international conditions are favorable for agreements (first and third rows in Figure 3.4). In these cases, international cooperation on geoengineering would better serve these other nations.

No Norms generally works best for the conflict measure because it is the U.S. policy that makes it least likely that a geoengineering system will ever be deployed. No Norms has high regret for three of the four measures when geoengineering is likely to succeed. But this policy

generally performs well at reducing conflict whenever international conditions are unfavorable for agreement (second, fourth, sixth, and eighth rows in Figure 3.4). This result owes to our assumption that No Norms can lead to conflicting and thus failed geoengineering tests. This, in turn, leads to fewer deployments even under catastrophic climate change. Without geoengineering deployments to blame, there is less opportunity for acrimony among nations suffering severe impacts from climate change. The same dynamic occurs, as described in “Evaluating Policies Across Multiple Scenarios” in Chapter Three, in the scenario in which international conditions are favorable for agreement and geoengineering has a significant impact on efforts to reduce emissions (first row in Figure 3.4). In addition, No Norms has low regret for the Collateral Damage measure in those cases in which geoengineering has little impact on efforts to reduce emissions and international conditions are unfavorable for agreement (fourth and eighth rows in Figure 3.4) because it decreases the chances that the United States and other nations will use geoengineering to protect themselves while abandoning weaker nations and natural ecosystems to suffer the impacts of climate change.

These patterns are summarized in Table 4.1. There are exceptions to these general patterns—that is, specific cases in which a policy could perform well (low regret) amid scenarios in which it generally fails.

For instance, Strong Norms generally has high regret when geoengineering is unlikely to succeed. But this is not the case when climate change proves catastrophic and geoengineering tests have a significant influence on efforts to reduce emissions (ninth and tenth rows in Figure 3.4). In this unpleasant scenario, no policy choice can prevent large impacts on the United States, large impacts on others, and significant collateral damage. In these scenarios and for these measures, Strong Norms has low regret, not because it performs well on any absolute scale but because all options are similarly bad. Note, however, that in these cases Strong Norms has high regret for the conflict measure because, relative to No Norms, it leads to more failed geoengineering deployments and thus more opportunity for acrimony among the nations damaged by these deployments.

Similarly, there are specific cases in which a policy performs poorly (high regret) amid scenarios in which it generally performs well. For instance, although Strong Norms generally has low regret when geoengineering is likely to succeed, it can have high regret for the Impacts to Others measure when international conditions are unfavorable for agreements. Under such conditions (second, fourth, sixth, and eighth rows in Figure 3.4), nations are most likely to conduct multilateral tests of geoengineering systems but are unlikely to reach any international agreement on deployment. The United States might thus deploy a geoengineering system unilaterally. The other nations would be better served by the Ban policy, under which they might have tested their own geoengineering systems. Strong Norms also has high regret for the conflict measure when international conditions are unfavorable for agreements because the No Norms policy would reduce the likelihood of a geoengineering deployment.

No Norms has low regret in all the scenarios considered here in which geoengineering is unlikely and climate change proves catastrophic. But this policy has high regret when geoengineering is unlikely to succeed, climate change proves severe, and the international conditions for agreement are favorable. Under such conditions (13th and 15th rows in Figure 3.4), significant reductions in greenhouse gas emissions are the only available path to avoiding climate impacts, and the Ban policy proves more effective at encouraging such reductions.

The Ban policy generally has low regret in all the scenarios considered here in which geoengineering is unlikely and climate change proves severe. But this policy can have high regret

Table 4.1
Summary of Vulnerabilities of Alternative Policies

Strategy	Vulnerable When				Because
	Geoengineering Potential	Climate	Geoengineering Influence	Political	
Strong Norms	Unlikely	Catastrophic	Weak	Favorable	The Strong Norms policy results in more successful tests but more failed deployments than does the No Norms policy, which permits conflicting and therefore failed tests.
Strong Norms	Unlikely	Severe	Strong	Favorable	The Strong Norms policy lead to weaker greenhouse gas reduction than does the Ban policy, while geoengineering does not help.
Strong Norms	Unlikely	Severe	Weak	Unfavorable	The Strong Norms policy results in more successful tests but more failed deployments than does the No Norms policy, which permits conflicting and therefore failed tests.
Ban	Likely	Catastrophic and severe ^a	Strong or weak	Favorable	A Ban policy increases the chances of geoengineering failure from inadequate testing, relative to the Strong Norms policy.
Ban	Unlikely	Catastrophic	Weak	Favorable	A Ban policy allows unilateral tests that can lead to failed deployments. The No Norms policy makes unilateral tests more difficult.
No Norms	Likely	Catastrophic	Strong or weak	Favorable	No Norms increases chances of geoengineering failure from inadequate testing, relative to the Strong Norms policy.
No Norms	Likely	Severe	Strong	Favorable	The No Norms policy leads to competing and failed experiments relative to the Strong Norms policy and therefore failed deployments.
No Norms	Unlikely	Severe	Strong	Favorable	Under these conditions, only significant emissions can prevent climate damage, and such reductions are most likely under a Ban policy.

^a In severe climate, there is low regret for impact on others.

when geoengineering tests have weak impact on efforts to reduce emissions and international conditions for agreement are unfavorable. Under such conditions (16th row in Figure 3.4), No Norms reduces the incidence of failed deployments. For the same reason, the Ban policy has high regret for the conflict measure with severe climate change whenever international conditions for agreement are unfavorable (14th and 16th rows in Figure 3.4). When geoengineering is unlikely and climate change proves catastrophic, the Ban policy always has high regret for the conflict measure and has high regret for all measures when geoengineering tests have little influence on efforts to reduce emissions. Under such conditions (11th and 12th rows in Figure 3.4), No Norms reduces the risk of failed geoengineering deployments, while the Ban policy adds little to efforts to reduce emissions.

Conclusions and Next Steps

Geoengineering presents a risky proposition. Deployment of an SRM geoengineering system could cause more harm than good. Yet, in some situations of extreme climate change, not deploying such a system could prove catastrophic. Today's policy choices involve questions about whether and how to conduct research that, in the future, might lead to a geoengineering deployment, not the deployment question itself. Yet, the ability to conduct a proper risk assessment of alternative proposals for the governance of geoengineering research has been limited because so much is deeply uncertain. Policymakers know very little about how the technology might work and, at least as importantly, know very little about how today's decisions about research might affect tomorrow's decisions about reducing greenhouse gas emissions and any geoengineering deployment.

This report demonstrates an approach to quantitative risk assessment under such conditions of deep uncertainty, based on a multiscenario vulnerability-and-response-option framework. The study uses a simple simulation model to project the consequences of three alternative U.S. government policies regarding geoengineering governance, contingent on uncertainties about the future potential of geoengineering technologies, the seriousness of climate change, the potential for international agreements limiting greenhouse gas emissions, and the impact that geoengineering research could have on such agreements. The model is run over many combinations of these uncertainties and the results used to identify potential vulnerabilities of each policy—that is, future conditions under which each would fail to meet U.S. government goals.

This report offers only the initial steps of a full RDM analysis. A full analysis could give decisionmakers confidence that they had thoroughly explored the vulnerabilities of a large set of alternative policy options and that could allow them to quantitatively weigh the trade-offs among these alternatives. The current study also makes many simplifying assumptions that a more-complete analysis would relax. Nonetheless, the work presented here suggests the potential of this approach for the systematic assessment of the risks of alternative geoengineering policies, suggests some of the key factors that a full analysis would need to include, and offers some initial suggestions toward the trade-offs the U.S. government might consider as it weighs whether and how to manage research on SRM geoengineering systems.

This analysis offers the following preliminary results for policymakers. If U.S. policymakers believe that some type of SRM technology is possible, they ought to prefer the Strong Norms policy. Under such conditions, this option outperforms the other two alternatives considered here because it increases the likelihood of a successful deployment in those cases in which geoengineering proves useful. It also reduces the likelihood of failed deployments by nations struggling to respond to serious climate impacts. But if U.S. policymakers believe that

no SRM geoengineering technology is likely to succeed, a Strong Norms policy could lead to overconfidence—enabling tests that, in turn, enable deployment of SRM systems that fail to work as expected. In addition, even if policymakers believe that SRM technology could succeed, under some conditions (unfavorable conditions for international agreements), the Strong Norms policy could advance some U.S. goals less successfully than a Ban or No Norms policy.

If they believe that successful SRM technology is unlikely, U.S. policymakers might prefer the Ban or No Norms option to Strong Norms. The Ban policy appears best if policymakers believe that climate change will likely prove severe but not catastrophic. Under such conditions, this option reduces the risks of overconfidence—deploying a geoengineering system that passes its tests but fails in practice. This option also increases the likelihood of reaching an international agreement to reduce greenhouse gas emissions.

In contrast, U.S. policymakers might prefer No Norms to Strong Norms or a Ban policy if they believe that climate change might prove catastrophic. Given our assumption that uncoordinated, conflicting geoengineering field tests can undermine all nations' ability to learn from such tests, the No Norms policy could ultimately prove more effective than a Ban policy at preventing geoengineering deployments that are unlikely to succeed and that might make matters worse under dire circumstances.

The analysis on which these observations are based, of course, would require significant refinement before it could move beyond a rough scan of vulnerabilities and systematically evaluate the trade-offs among alternative policy options.

First, the existing simulation could be used to conduct a more-complete RDM analysis of the alternative policy options. The present simulation contains many more uncertain input parameters than those varied in the analysis presented here. In particular, the conditional probability tables for the 17 nodes in the Bayesian network derive from 93 independent parameters. A full RDM analysis would replace the 16 scenarios in Figure 3.4 with a database of several thousand cases using a statistical experiment design (e.g., Latin hypercube sample; see Saltelli, Chan, and Scott, 2000) spanning all the parameters defining the values in the conditional probability tables. To replace Figure 3.4 and Table 4.1, the analysis would then summarize these thousands of cases using scenario-discovery algorithms that identify easy-to-interpret clusters of cases that best represent the vulnerabilities of each policy (Bryant and Lempert, 2010; Groves and Lempert, 2007). This scenario discovery would provide information about the importance of the four uncertain parameters considered in this study compared to other key uncertainties, as well as a more-exhaustive exploration of the policies' potential vulnerabilities, including some that decisionmakers might regard as surprises (Lempert, 2007; Lempert, Popper, and Bankes, 2002). Once the RDM analysis identifies the policies' vulnerable scenarios, it could then suggest how likely these scenarios would need to be in order for policymakers to choose one policy option over another. For instance, such an analysis would likely highlight the importance of our assumption that uncoordinated geoengineering tests can undermine the ability to learn from such tests to the potential vulnerabilities of the Ban policy. An enhanced analysis might also suggest how much confidence policymakers would need to have in geoengineering technologies in order to justify choosing Strong Norms over the Ban policy.

The data used to inform the current simulation could also be significantly improved. At present, the ranges and nominal values in the conditional probability tables represent only our judgments. Analysts could conduct a formal elicitation process of experts in such fields as geoengineering technology and environmental politics to inform the range of values for these

parameters (Morgan et al., 2009). Combined with the RDM uncertainty analysis, these expert elicitations would provide the best available information to evaluate trade-offs among the decision options.

Finally, enhancements to the simulation could expand the range of policy options considered and more-thoroughly explore each strategy's vulnerabilities. This would improve the study's ability to suggest more-robust strategies. The current simulation has only a very crude time dimension. Increasing the model's temporal resolution could permit evaluation of policies that evolve over time in response to new information. For instance, policy might exploit a failed geoengineering test to encourage a more-urgent effort to reduce greenhouse gases. A model that differentiates between laboratory and field experiments would allow the consideration of near-term policies intermediate between Strong Norms and a Ban policy, ones that could pursue aggressive laboratory and computer testing before determining the conditions for testing SRM in the field. It is not immediately obvious how considering more-adaptive policies might affect the results presented here. For instance, an adaptive Strong Norms policy might prove less vulnerable because it could respond appropriately to initial research that might suggest that geoengineering technology is likely to fail. However, such research might fail to alleviate a Strong Norms policy's potential overconfidence.

A full RDM analysis using such an enhanced simulation could help examine such questions. Such an analysis would enable policymakers to explore ways to reduce the potential for any overconfidence in Strong Norms and identify the scenarios in which it might prove most dangerous. Such an analysis could help design more-robust policies than considered here by iterating through several cycles of vulnerability-and-response-option analyses, at each stage evaluating actions to reduce vulnerabilities of policies like Strong Norms, and then exploring for any new vulnerabilities introduced by those actions. (See Lempert, Groves, et al., 2006, for an example of such an iterative RDM analysis.)

Disaggregating other nodes could also enrich the evaluation of policies. For instance, a revised model might consider separately the international conditions affecting prospects for agreements on emission reductions and on geoengineering deployments. Such disaggregation might also allow consideration of near-term policy choices that could affect the types of future agreements reached. For instance, the venues used by the U.S. government to pursue or block norms on geoengineering research might affect the potential for international agreements on greenhouse gas reductions. Appropriate risk-sharing and conflict-resolution mechanisms consistent with Victor's concept of cocktail geoengineering might, for instance, reduce the vulnerabilities the current Strong Norms policy exhibits for the conflict measure and the impacts on other nations.

Relaxing the state-centered focus of the current analysis might also enrich understanding of the "stickiness" of alternative near-term policies—that is, help evaluate how initial policy choices might create political coalitions that would make some future policy adjustments easier to pursue than others. The formation of such coalitions can play an important role in the longevity of various policies (Patashnik, 2003). For instance, some research policies might encourage private-sector actors to begin geoengineering programs that might constrain future government policy.

Despite the potential usefulness of such enhancements, the current study provides useful contributions to the growing policy debates over geoengineering. It demonstrates an approach for a risk analysis that can address questions regarding the governance of geoengineering research despite the deep uncertainties involved. It identifies some of the steps needed to con-

duct a more-complete assessment of such questions. Finally, it provides some initial policy results. In particular, the Royal Society and other commentators have recommended policies similar to Strong Norms as the best means for managing geoengineering activities in the near term. The analysis presented here suggests that such recommendations might prove reasonable but rest on several key assumptions that might not necessarily hold. Strong Norms assumes that there is a reasonable chance that SRM technology could prove successful. If such technology is unlikely to work, the Strong Norms policy might increase the likelihood of an unsuccessful deployment in response to abrupt climate change. In addition, Strong Norms assumes that international conditions are reasonably favorable for agreements governing greenhouse gas emission reductions and geoengineering deployments. If such conditions do not hold, the Strong Norms policy might increase the likelihood of a conflict-provoking, unilateral geoengineering deployment.

Whether or not they share these assumptions, understanding the potential vulnerabilities of the Strong Norms and other options can help policymakers understand the risks posed by alternative decisions, help them design potentially more-robust options, and better manage the risks and opportunities presented by geoengineering technology.

Types of Geoengineering Technologies

The literature describes two general types of geoengineering technologies: CDR and SRM. This appendix describes key representatives of each type, although the study focuses on the latter.

Carbon-removal technologies limit greenhouse gas-induced climate change by extracting CO₂ directly from the atmosphere and then storing it in some type of reservoir.¹ CDR approaches can include reforestation projects, in which growing trees extract and store carbon. Although such reforestation projects are low cost and might provide many environmental benefits related to the preservation of habitat, they offer only limited capacity to store carbon compared to the magnitude of emissions from fossil-fuel combustion. CDR approaches also include ocean-iron or ocean-nitrogen fertilization, in which iron compounds are added to selected regions of the ocean to stimulate growth of phytoplankton that absorbs CO₂ and transfers it to the deep ocean when the organisms die. This approach has been much discussed, though recent ocean-iron tests have not initially proved promising.

The natural carbon cycle removes CO₂ from the atmosphere on a long-term basis through the weathering of carbonate and silicate rocks. Some CDR approaches aim to accelerate these natural processes (House et al., 2007). For instance, one proposal would mix finely ground silicate materials to agricultural soils, which would then absorb CO₂ from the atmosphere. Estimates suggest that adding an amount of silicate equal to about twice the current global rate of coal mining to soils could counteract all current anthropogenic greenhouse gas emissions. Although thus potentially very effective, such approaches would likely prove costly, and their environmental impacts on soils, although not obviously adverse, are poorly understood. Another approach with potentially unlimited capacity involves devices that capture CO₂ directly from the air for subsequent sequestration in geological formations. For instance, large radiator-like screens could be deployed that would use highly alkaline solutions to absorb the CO₂ from the air passing through. Such air-capture systems might prove significantly more costly than extracting CO₂ from the far denser streams in the effluents of power plants. But the economics of air capture might be aided by the flexibility to site facilities near storage sites or near sources of excess energy supply. The environmental impacts of such systems involve any danger of leakage from storage sites and the need to occupy large tracts of land.

The characteristics of most CDR approaches suggest that they would fit relatively easily into proposed institutions for addressing climate change, would pose few new security issues, and could largely be managed by national governments. Overall, CDR approaches act slowly. None could change the atmospheric concentration of greenhouse gases by more than a frac-

¹ This section draws from the technology summaries in Royal Society (2009).

tion of a percent per year and thus, at best, would take decades to counteract even the current buildup of such gases. Although a single or few nations could, in principle, implement CDR approaches unilaterally, the costs, at best, will not likely prove much smaller than many approaches to limit greenhouse gas emissions. Thus, CDR would most likely be employed as one part of an international effort to limit atmospheric concentrations of greenhouse gases. For example, such an effort might finance CDR systems by allowing them to serve as offsets or sell permits in a cap-and-trade system. Many reforestation projects have already been funded in this fashion using the clean-development mechanism of the Kyoto Protocol.²

Direct carbon-removal systems do pose governance issues. Any approach that generates offsets or permits would require standards for monitoring, validation, and, perhaps, discussions of how much ultimately needs to be removed. They would require arrangements to address liability for leakage from storage decades into the future. Any adverse environmental impacts of CDR approaches might prove relatively minor, and any that do occur should prove relatively localized in the land or ocean area surrounding the implementation. Nonetheless, tests of CDR systems will require appropriate regulation and safeguards. For instance, recent negotiations under the London ocean dumping convention have established rules (International Maritime Organization, 2008) that govern disclosure and impact assessments of ocean-fertilization tests. But most CDR systems will fall under national jurisdictions, only minimally crossing national boundaries, so any environmental impacts could be managed appropriately by existing national institutions.

In contrast, SRM approaches would prove much more difficult to manage. SRM technologies reflect a small amount of incoming solar radiation back into space, thereby cooling the planet to offset some of the warming from increased concentrations of greenhouse gases. One SRM approach that increases the reflectivity of the earth's surface by brightening human infrastructure (for instance, by painting roofs white) has few adverse side effects but could counteract only a small fraction of warming.

Other SRM approaches offer much more-significant impacts but with more-significant risks. One approach would inject sulfate particles into the lower stratosphere in order to increase its reflectivity. The effects would mimic those from large volcanic eruptions, such as Mount Pinatubo in 1991, which cooled the earth by about 0.5 degree Celsius for about two years. Calculations suggest that reducing solar input by about 2 percent would counteract the warming from a doubling of greenhouse gas concentrations in the atmosphere. This reduction could require injecting about 1 million to 5 million tons of sulfur, perhaps sulfate aerosols, annually to an altitude of about 20 kilometers. This amount is equivalent to a small portion of the annual global air transport.³ The economics are surprisingly feasible, on the order of tens of billions of dollars, one or more orders of magnitude less than the estimates of costs of reducing emissions of greenhouse gases. Climate model calculations and scientific observations of the aftermath of the Mount Pinatubo eruption suggest that a future climate that combines a doubling of greenhouse gas concentrations with sulfate aerosol geoengineering would be much closer to the current climate than one with a doubling of greenhouse gas concentrations

² Kyoto Protocol to the United Nations Framework Convention on Climate Change. The clean-development mechanism is defined in Article 12 of the convention.

³ Each of the major world airports handles several million tons of cargo per year (see Airports Council International, 2009). The Berlin airlift of 1948–1949 moved comparable amounts of material by air. Depending on the required deposition altitude, the transport issue might become more significant.

alone. But studies also suggest that Mount Pinatubo's eruption led to substantial decreases in precipitation over land areas and about a 2-percent reduction of stratospheric ozone levels. Initial model simulations suggest that sulfate geoengineering could similarly reduce precipitation from the Asian monsoons, cause other adverse changes in regional weather patterns, and thin the ozone layer. Sulfate aerosols also do not address other impacts of increased atmospheric concentrations of CO_2 , such as ocean acidification.

Another SRM approach would enhance cloud albedo by spraying fine particulates into the low-level marine clouds that cover about one-quarter of the earth's ocean surface. Such particles would serve as condensation nuclei, whitening the clouds and cooling the earth by increasing its reflectivity. Solar-reflector approaches would place sun shields in space, either in near-earth orbit or at the Lagrange (L1) point about 1.5 million km from earth where the gravitational attraction of the earth and sun balance. Similarly to the stratospheric aerosol approaches, the effectiveness and potential adverse effects of both these approaches are poorly understood. Cloud-albedo enhancement could prove similarly cost-effective as stratospheric aerosols—estimates suggest that a fleet of roughly 1,500 ships could implement the cloud-whitening approach—but could alter weather patterns and ocean currents. Deploying space-based reflectors could prove significantly more expensive and could also affect the climate in different regions differently.

Overall, several SRM technologies—including sulfate aerosols, cloud whitening, and near-earth orbit solar reflectors—provide a quick-acting, potentially reversible, and possibly inexpensive approach to countering some of the most-important impacts of climate change. However, these technologies also present significant risks of adverse unintended consequences, most of which are, at present, very poorly understood.

Relationships

This appendix summarizes the current behavior of the eight relationship nodes in Figures 3.1, 3.2, and 3.3. This behavior is chosen to illustrate the complex relationship and large dynamic range of possible outcomes. The conditional probability tables and a description of how they were created are given in Appendix C. Future work will vary the conditional probability tables and therefore increase the range of possible outcomes described in this appendix. Table B.1

Table B.1
Relationships Among Nodes in Bayesian Network

Node	State	Influences	Is Influenced by
EmissionsReductions	Strong, Weak	ClimateImpactsUS ^a ClimateImpactOthers ^a CollateralDamage ^a USDeploys OthersDeploy	PotentialForAgreement ^b GeoEngineeringInfluence ^b USTests OthersTest DeploymentAgreement
USTests	Multilateral, Unilateral, NoTest	TestResults EmissionsReductions	Norms ^c
OthersTest	Unilateral, NoTest	TestResults EmissionsReductions	Norms ^c
TestResults	GlobalSuccess, SuccessforUS, SuccessforOthers, Failure, NoResults	USDeploys OthersDeploy DeploymentResults	GeoEngineeringPotential ^b USTests OthersTest
DeploymentAgreement	GoverningRules, Ban, NoAgreement	USDeploys OthersDeploy EmissionsReductions	Norms ^c PotentialforAgreement ^b
USDeploys	Multilateral, Unilateral, NoDeployment	DeploymentResults	ClimateChange ^b EmissionsReductions DeploymentAgreement TestResults
OthersDeploy	Unilateral, NoDeployment	DeploymentResults	ClimateChange ^b EmissionsReductions DeploymentAgreement TestResults
DeploymentResults	GlobalSuccess, SuccessforUS, SuccessforOthers Failure NoResults	ClimateImpactsUS ^a ClimateImpactsOthers ^a CollateralDamage ^a Conflict ^a	GeoEngineeringPotential ^b TestResults USDeploys OthersDeploy

^a Measure.

^b Uncertainty.

^c Lever.

summarizes each of the eight relationship nodes, their states, and the nodes that each node influences and is influenced by.

EmissionsReductions Node

This node is influenced by international political conditions and whether or not undertaking geoengineering research affects the international will to cut greenhouse gases. The latter is governed by the scenario choice of either a strong or weak geoengineering influence. The Weak condition dissociates geoengineering tests and decisions from progress on global greenhouse gas–reduction progress.

Under a strong influence scenario, anything that indicates that geoengineering is a serious option leads to weak reductions in global emissions. This could be a result of serious international negotiations (shown by governing rules) or decisions to conduct large-scale, expensive field tests of geoengineering. These tests might take decades to complete and could lead to decisions to wait for the results before instituting emission reductions that have adverse economic impacts. An international agreement on how to deploy geoengineering creates a belief that collateral damage will be controlled and geoengineering is possible. Of the three choices on deployment agreements, governing rules are likely to have the greatest negative impact on emission reductions because formal agreements signal that the global community believes that geoengineering might well be needed and might work. Thus, there is even a small possibility that strong agreements, even without testing, can lead to weak emission reductions because the community believes that, in a crisis, geoengineering might be deployed with no prior testing.

We assume that no agreement suggests that geoengineering is less likely to be pursued and therefore has less of an impact on greenhouse gas reductions, especially if there are no tests conducted.

Large-scale greenhouse gas reduction will require a global effort; therefore, unfavorable political and economic conditions lead to weak greenhouse gas reduction.

Test Decision Nodes

The tests considered here are large-scale field tests. In actuality, these tests will probably be a series of experimental campaigns lasting perhaps decades. A decision to test geoengineering might be made either unilaterally or multilaterally. A multilateral test is conducted by a group of nations working collaboratively. Such tests are likely to give information about both potential global and regional impacts. Multilateral tests are designed to maximize favorable global effects. In this network, it is assumed that any multilateral testing will include the United States. A unilateral test is conducted by a single nation or a small group of nations acting alone. Such tests are most likely to be designed to yield favorable regional impacts rather than global impacts. The decision to test is influenced by three possible levels of norms. Strong Norms policies favor multilateral tests, while No Norms policies favor unilateral tests. A norm that promotes a ban restricts testing by the United States, but less so by others due to the assumed asymmetry in meeting one's obligations to international agreements discussed earlier. Only under No Norms are multiple tests possible as the United States abides by a ban and, under Strong Norms, tests are coordinated to prevent interference. Others' tests might take place

even in the face of a ban if countries feel they need an option in the event of catastrophic climate change.

For simplicity, we have chosen to omit laboratory tests and computer simulations, although they would certainly precede field testing. They might have an impact on greenhouse gas reduction if geoengineering influence were high and certainly need to be considered in a more-extensive model. We believe it unlikely that either the simulations or laboratory tests would be negative enough to definitively preclude any field testing.

TestResults Node

Test success must be defined by success compared to what and for whom. Success means that the climate is considerably better with geoengineering than without. It does not mean a complete lack of droughts, desertification, or other effects. A global success means that there are widespread beneficial impacts. Success for others or success for the United States means that the specific political entity is favorably affected, but many other areas suffer adverse impacts.

Failure means that there is no indication that the proposed technology provides a beneficial impact and, in fact, might actually make the situation worse. No results indicate that no test was conducted.

The choice of norm policy affects only the probability that one conducts the tests, not the relative outcome of the tests.

Tests might yield results other than what was expected or desired. For example, a multilateral test could show that there are no globally beneficial results but that there are significant benefits to others. A unilateral U.S. test could show beneficial effects for others. Others might then base future decisions on these results.

The test results are affected by the uncertainty in the technology and who conducts the tests. Multilateral tests are more likely to have beneficial results than unilateral tests. If multiple unilateral tests are conducted, the likelihood of failure is high due to the difficulty of interpreting results in the presence of multiple drivers and competing climatic effects. Multiple tests are highly unlikely under a Strong Norms policy. An assessment that geoengineering is unlikely to succeed (an X variable) leads only to a reduced likelihood of success, not a zero probability.

DeploymentAgreement Node

Similar to norms, deployment agreements have three states: GoverningRules, Ban, and NoAgreement. A GoverningRules state implies that geoengineering is taken seriously by the world community. In this state, governments reach an international agreement governing the deployment of any geoengineering system. The agreement might include rules and funding for providing compensation for any adverse regional impacts. The agreement might also include requirements that any geoengineering deployment includes cocktail technologies to reduce collateral impacts.

A Ban state indicates an international agreement to prohibit any deployment of geoengineering. Monitoring protocols and international sanctions for countries that deploy outside of the agreement might also be included.

The likelihood of any of these states occurring is closely correlated to the existing norms, moderated by whether or not the global community is cooperating.

Deployment-Decision Nodes: USDeploys and OthersDeploy

The deployment decision will depend on the presence or absence of test results, the urgency of the climate situation as determined by progress on greenhouse gas reduction, climate predictions, and a general belief in the possibility of geoengineering.

Under catastrophic conditions, the only instance in which geoengineering is not deployed is as a result of a failed test. A lack of testing does not halt deployment. Neither emission reduction nor international agreement halts deployment. This could result in multiple geoengineering strategies being implemented simultaneously. If there is a global test success, the United States will deploy as part of a multilateral team. If there is only a U.S. success, the United States will still deploy unilaterally in spite of any governing agreements. The United States does not deploy if there is a success for others and no U.S. success. Others' behavior is similar—others deploy except in the presence of a failure or U.S. success.

Under severe but not catastrophic climate conditions, deployment decisions are less likely. The United States will live up to its international commitments. The United States is likely to deploy only after successful global tests and in the presence of weak emission reduction. Even with success for the United States, weak emission reductions, and severe climate predictions, the United States obeys its international agreements and does not deploy unilaterally if there are governing rules. This is the negative aspect of strong governance.

Others' behavior mirrors U.S. behavior for catastrophic climate conditions but not severe conditions. It is assumed that others might deploy in severe conditions in spite of either strong agreements or bans. This reflects the asymmetry between U.S. behavior and others' behavior.

If the climatic conditions are mild, geoengineering is not deployed under any circumstances.

DeploymentResults Node

Success is assumed to be likely if geoengineering potential is assessed as possible, but success is reduced by either a failed test or no test. Successful results are also less likely if the deployment occurs contrary to test results—for example, if others deploy based on successful U.S. tests. The success rate is cut substantially if multiple deployments are undertaken unless one of the deployments is multilateral. U.S. and others' success cannot happen simultaneously—that is, only a global success can have low impact on both the United States and others. If geoengineering is judged to be impossible, the success rate is cut but not reduced to zero.

Conditional Probability Tables

The conditional probability tables were constructed to demonstrate the complex relationships between current policy choices and the eventual availability and efficacy of geoengineering. Using our judgment for each of the “influenced-by” nodes given in Table B.1 in Appendix B, we generated relationships that detail the strength of that influence on a particular node in the chain (i.e., how important that factor is to the probability of achieving a particular outcome).

This process is illustrated in Table C.1 for the EmissionsReductions node. This node gives the probability that strong greenhouse gas reductions will occur. The EmissionsReductions node is influenced by five factors: geoengineering influence, potential for agreements, deployment agreement, U.S. test, and others’ tests. As illustrated in Figure C.1, we initially determined that weak geoengineering influence would have no impact on greenhouse gas reductions (0 percent in the table), while unfavorable potential for agreements would reduce the likelihood of greenhouse gas reduction by 50 percent. The influence of each condition is shown in the grayed-out portion of Table C.1, dependent on the condition of each of the five parent nodes. The final probabilities for the strong state of the EmissionsReductions node are simply computed as the product of the influences; weak is then computed as the quantity (1 - strong). These probabilities are bolded in Table C.1.

The advantage of using these influence tables is that the influences can be varied over wide ranges in a consistent and rapid manner to determine threshold behavior and the most-sensitive critical factors. For the purposes of this report, single values were chosen with the intent of demonstrating the wide dynamic range of the results.

The actual probability tables are given in Tables C.2 through C.13. The probabilities for the states of each node are given on the left as a function of the states of the parent nodes on the right.

Table C.1
Emission-Reduction Influence Factors

GeoEngineering Influence (Strong, Weak)	PotentialFor Agreements (Favorable, Unfavorable)	Deployment Agreement (Rules, Ban, None)	US Tests (Multi, Uni, None)	Others Test (Uni, None)
100%	100%	50%	50%	50%
0%	50%	100%	100%	100%
		100%	100%	

GE Influence	PotentialFor Agreements (Favorable, Unfavorable)	Deployment Agreement	US Tests	Others Test	Strong	Weak
Strong	Favorable	Rules	Multilateral	Unilateral	0.1250	0.8750
Strong	Favorable	Rules	Multilateral	None	0.2500	0.7500
Strong	Favorable	Rules	Unilateral	Unilateral	0.1250	0.8750
Strong	Favorable	Rules	Unilateral	None	0.2500	0.7500
Strong	Favorable	Rules	None	Unilateral	0.2500	0.7500
Strong	Favorable	Rules	None	None	0.5000	0.5000
Strong	Favorable	Ban	Multilateral	Unilateral	0.2500	0.7500
Strong	Favorable	Ban	Multilateral	None	0.5000	0.5000
Strong	Favorable	Ban	Unilateral	Unilateral	0.2500	0.7500
Strong	Favorable	Ban	Unilateral	None	0.5000	0.5000
Strong	Favorable	Ban	None	Unilateral	0.5000	0.5000
Strong	Favorable	Ban	None	None	1.0000	0.0000
Strong	Favorable	None	Multilateral	Unilateral	0.2500	0.7500
Strong	Favorable	None	Multilateral	None	0.5000	0.5000
Strong	Favorable	None	Unilateral	Unilateral	0.2500	0.7500
Strong	Favorable	None	Unilateral	None	0.5000	0.5000
Strong	Favorable	None	None	Unilateral	0.5000	0.5000
Strong	Favorable	None	None	None	1.0000	0.0000
Strong	Unfavorable	Rules	Multilateral	Unilateral	0.0625	0.9375
Strong	Unfavorable	Rules	Multilateral	None	0.1250	0.8750



GE Influence	PotentialFor Agreements	Deployment Agreement	US Tests	Weak
100%	100%	50%	50%	50%
100%	100%	50%	50%	100%
100%	100%	50%	50%	50%
100%	100%	50%	50%	100%
100%	100%	50%	100%	50%
100%	100%	50%	100%	100%
100%	100%	100%	50%	50%
100%	100%	100%	50%	100%
100%	100%	100%	50%	50%
100%	100%	100%	50%	100%
100%	100%	100%	100%	50%
100%	100%	100%	100%	100%
100%	100%	100%	100%	100%
100%	100%	100%	50%	50%
100%	100%	100%	50%	100%
100%	100%	100%	50%	100%
100%	100%	100%	100%	50%
100%	100%	100%	100%	100%
100%	50%	50%	50%	50%
100%	50%	50%	50%	100%

Table C.2
Conditional Probability Table: DeploymentAgreement

GoverningRules	Ban	NoAgreement	Norms	PotentialForAgreements
0.9	0.1	0	Strong	Favorable
0.25	0.125	0.625	Strong	UnFavorable
0.1	0.9	0	Ban	Favorable
0.125	0.25	0.625	Ban	UnFavorable
0.09	0.09	0.82	NoNorm	Favorable
0.025	0.025	0.95	NoNorm	UnFavorable

Table C.3
Conditional Probability Table: EmissionsReductions

Strong	Weak	GeoEngineering Influence	PotentialFor Agreements	Deployment Agreement	USTests	OthersTest
0.125	0.875	Strong	Favorable	GoverningRules	Multilateral	Unilateral
0.25	0.75	Strong	Favorable	GoverningRules	Multilateral	NoTest
0.125	0.875	Strong	Favorable	GoverningRules	Unilateral	Unilateral
0.25	0.75	Strong	Favorable	GoverningRules	Unilateral	NoTest
0.25	0.75	Strong	Favorable	GoverningRules	NoTest	Unilateral
0.5	0.5	Strong	Favorable	GoverningRules	NoTest	NoTest
0.25	0.75	Strong	Favorable	Ban	Multilateral	Unilateral
0.5	0.5	Strong	Favorable	Ban	Multilateral	NoTest
0.25	0.75	Strong	Favorable	Ban	Unilateral	Unilateral
0.5	0.5	Strong	Favorable	Ban	Unilateral	NoTest
0.5	0.5	Strong	Favorable	Ban	NoTest	Unilateral
1	0	Strong	Favorable	Ban	NoTest	NoTest
0.25	0.75	Strong	Favorable	NoAgreement	Multilateral	Unilateral
0.5	0.5	Strong	Favorable	NoAgreement	Multilateral	NoTest
0.25	0.75	Strong	Favorable	NoAgreement	Unilateral	Unilateral
0.5	0.5	Strong	Favorable	NoAgreement	Unilateral	NoTest
0.5	0.5	Strong	Favorable	NoAgreement	NoTest	Unilateral
1	0	Strong	Favorable	NoAgreement	NoTest	NoTest
0.0625	0.9375	Strong	UnFavorable	GoverningRules	Multilateral	Unilateral
0.125	0.875	Strong	UnFavorable	GoverningRules	Multilateral	NoTest
0.0625	0.9375	Strong	UnFavorable	GoverningRules	Unilateral	Unilateral
0.125	0.875	Strong	UnFavorable	GoverningRules	Unilateral	NoTest
0.125	0.875	Strong	UnFavorable	GoverningRules	NoTest	Unilateral

Table C.3—Continued

Strong	Weak	GeoEngineering Influence	PotentialFor Agreements	Deployment Agreement	USTests	OthersTest
0.25	0.75	Strong	UnFavorable	GoverningRules	NoTest	NoTest
0.125	0.875	Strong	UnFavorable	Ban	Multilateral	Unilateral
0.25	0.75	Strong	UnFavorable	Ban	Multilateral	NoTest
0.125	0.875	Strong	UnFavorable	Ban	Unilateral	Unilateral
0.25	0.75	Strong	UnFavorable	Ban	Unilateral	NoTest
0.25	0.75	Strong	UnFavorable	Ban	NoTest	Unilateral
0.5	0.5	Strong	UnFavorable	Ban	NoTest	NoTest
0.125	0.875	Strong	UnFavorable	NoAgreement	Multilateral	Unilateral
0.25	0.75	Strong	UnFavorable	NoAgreement	Multilateral	NoTest
0.125	0.875	Strong	UnFavorable	NoAgreement	Unilateral	Unilateral
0.25	0.75	Strong	UnFavorable	NoAgreement	Unilateral	NoTest
0.25	0.75	Strong	UnFavorable	NoAgreement	NoTest	Unilateral
0.5	0.5	Strong	UnFavorable	NoAgreement	NoTest	NoTest
1	0	Weak	Favorable	GoverningRules	Multilateral	Unilateral
1	0	Weak	Favorable	GoverningRules	Multilateral	NoTest
1	0	Weak	Favorable	GoverningRules	Unilateral	Unilateral
1	0	Weak	Favorable	GoverningRules	Unilateral	NoTest
1	0	Weak	Favorable	GoverningRules	NoTest	Unilateral
1	0	Weak	Favorable	GoverningRules	NoTest	NoTest
1	0	Weak	Favorable	Ban	Multilateral	Unilateral
1	0	Weak	Favorable	Ban	Multilateral	NoTest
1	0	Weak	Favorable	Ban	Unilateral	Unilateral
1	0	Weak	Favorable	Ban	Unilateral	NoTest
1	0	Weak	Favorable	Ban	NoTest	Unilateral
1	0	Weak	Favorable	Ban	NoTest	NoTest
1	0	Weak	Favorable	NoAgreement	Multilateral	Unilateral
1	0	Weak	Favorable	NoAgreement	Multilateral	NoTest
1	0	Weak	Favorable	NoAgreement	Unilateral	Unilateral
1	0	Weak	Favorable	NoAgreement	Unilateral	NoTest
1	0	Weak	Favorable	NoAgreement	NoTest	Unilateral
1	0	Weak	Favorable	NoAgreement	NoTest	NoTest
0.5	0.5	Weak	UnFavorable	GoverningRules	Multilateral	Unilateral
0.5	0.5	Weak	UnFavorable	GoverningRules	Multilateral	NoTest

Table C.3—Continued

Strong	Weak	GeoEngineering Influence	PotentialFor Agreements	Deployment Agreement	USTests	OthersTest
0.5	0.5	Weak	UnFavorable	GoverningRules	Unilateral	Unilateral
0.5	0.5	Weak	UnFavorable	GoverningRules	Unilateral	NoTest
0.5	0.5	Weak	UnFavorable	GoverningRules	NoTest	Unilateral
0.5	0.5	Weak	UnFavorable	GoverningRules	NoTest	NoTest
0.5	0.5	Weak	UnFavorable	Ban	Multilateral	Unilateral
0.5	0.5	Weak	UnFavorable	Ban	Multilateral	NoTest
0.5	0.5	Weak	UnFavorable	Ban	Unilateral	Unilateral
0.5	0.5	Weak	UnFavorable	Ban	Unilateral	NoTest
0.5	0.5	Weak	UnFavorable	Ban	NoTest	Unilateral
0.5	0.5	Weak	UnFavorable	Ban	NoTest	NoTest
0.5	0.5	Weak	UnFavorable	NoAgreement	Multilateral	Unilateral
0.5	0.5	Weak	UnFavorable	NoAgreement	Multilateral	NoTest
0.5	0.5	Weak	UnFavorable	NoAgreement	Unilateral	Unilateral
0.5	0.5	Weak	UnFavorable	NoAgreement	Unilateral	NoTest
0.5	0.5	Weak	UnFavorable	NoAgreement	NoTest	Unilateral
0.5	0.5	Weak	UnFavorable	NoAgreement	NoTest	NoTest

Table C.4
Conditional Probability Table: TestResults

GlobalSuccess	SuccessforUS	Successfor Others	Failure	NoResults	GeoEngineering Potential	USTests	OthersTest
0.9	0.0125	0.0125	0.075	0	Likely	Multilateral	Unilateral
0.9	0.0783	0.0217	0	0	Likely	Multilateral	NoTest
0	0.125	0.125	0.75	0	Likely	Unilateral	Unilateral
0.1	0.7941	0.1059	0	0	Likely	Unilateral	NoTest
0.1	0.1059	0.7941	0	0	Likely	NoTest	Unilateral
0	0	0	0	1	Likely	NoTest	NoTest
0.45	0.0625	0.0625	0.425	0	Unlikely	Multilateral	Unilateral
0.45	0.3913	0.1087	0.05	0	Unlikely	Multilateral	NoTest
0	0.0625	0.0625	0.875	0	Unlikely	Unilateral	Unilateral
0.05	0.4412	0.0588	0.45	0	Unlikely	Unilateral	NoTest
0.05	0.0588	0.4412	0.45	0	Unlikely	NoTest	Unilateral
0	0	0	0	1	Unlikely	NoTest	NoTest

Table C.5
Conditional Probability Table: OthersTest

Unilateral	NoTest	Norms
0	1	Strong
0.75	0.25	Ban
1	0	NoNorm

Table C.6
Conditional Probability Table: OthersDeploy

Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
1	0	Catastrophic	GoverningRules	Strong	GlobalSuccess
0	1	Catastrophic	GoverningRules	Strong	SuccessforUS
1	0	Catastrophic	GoverningRules	Strong	SuccessforOthers
0	1	Catastrophic	GoverningRules	Strong	Failure
1	0	Catastrophic	GoverningRules	Strong	NoResults
1	0	Catastrophic	GoverningRules	Weak	GlobalSuccess
0	1	Catastrophic	GoverningRules	Weak	SuccessforUS
1	0	Catastrophic	GoverningRules	Weak	SuccessforOthers
0	1	Catastrophic	GoverningRules	Weak	Failure
1	0	Catastrophic	GoverningRules	Weak	NoResults
1	0	Catastrophic	Ban	Strong	GlobalSuccess
0	1	Catastrophic	Ban	Strong	SuccessforUS
1	0	Catastrophic	Ban	Strong	SuccessforOthers
0	1	Catastrophic	Ban	Strong	Failure
1	0	Catastrophic	Ban	Strong	NoResults
1	0	Catastrophic	Ban	Weak	GlobalSuccess
0	1	Catastrophic	Ban	Weak	SuccessforUS
1	0	Catastrophic	Ban	Weak	SuccessforOthers
0	1	Catastrophic	Ban	Weak	Failure
1	0	Catastrophic	Ban	Weak	NoResults
1	0	Catastrophic	NoAgreement	Strong	GlobalSuccess
0	1	Catastrophic	NoAgreement	Strong	SuccessforUS
1	0	Catastrophic	NoAgreement	Strong	SuccessforOthers
0	1	Catastrophic	NoAgreement	Strong	Failure
1	0	Catastrophic	NoAgreement	Strong	NoResults
1	0	Catastrophic	NoAgreement	Weak	GlobalSuccess

Table C.6—Continued

Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
0	1	Catastrophic	NoAgreement	Weak	SuccessforUS
1	0	Catastrophic	NoAgreement	Weak	SuccessforOthers
0	1	Catastrophic	NoAgreement	Weak	Failure
1	0	Catastrophic	NoAgreement	Weak	NoResults
0.09375	0.90625	Severe	GoverningRules	Strong	GlobalSuccess
0	1	Severe	GoverningRules	Strong	SuccessforUS
0.1875	0.8125	Severe	GoverningRules	Strong	SuccessforOthers
0	1	Severe	GoverningRules	Strong	Failure
0.0468795	0.953121	Severe	GoverningRules	Strong	NoResults
0.375	0.625	Severe	GoverningRules	Weak	GlobalSuccess
0	1	Severe	GoverningRules	Weak	SuccessforUS
0.75	0.25	Severe	GoverningRules	Weak	SuccessforOthers
0	1	Severe	GoverningRules	Weak	Failure
0.1875	0.8125	Severe	GoverningRules	Weak	NoResults
0.09375	0.90625	Severe	Ban	Strong	GlobalSuccess
0	1	Severe	Ban	Strong	SuccessforUS
0.1875	0.8125	Severe	Ban	Strong	SuccessforOthers
0	1	Severe	Ban	Strong	Failure
0.0468795	0.953121	Severe	Ban	Strong	NoResults
0.375	0.625	Severe	Ban	Weak	GlobalSuccess
0	1	Severe	Ban	Weak	SuccessforUS
0.75	0.25	Severe	Ban	Weak	SuccessforOthers
0	1	Severe	Ban	Weak	Failure
0.1875	0.8125	Severe	Ban	Weak	NoResults
0.125	0.875	Severe	NoAgreement	Strong	GlobalSuccess
0	1	Severe	NoAgreement	Strong	SuccessforUS
0.25	0.75	Severe	NoAgreement	Strong	SuccessforOthers
0	1	Severe	NoAgreement	Strong	Failure
0.0625	0.9375	Severe	NoAgreement	Strong	NoResults
0.5	0.5	Severe	NoAgreement	Weak	GlobalSuccess
0	1	Severe	NoAgreement	Weak	SuccessforUS
1	0	Severe	NoAgreement	Weak	SuccessforOthers
0	1	Severe	NoAgreement	Weak	Failure

Table C.6—Continued

Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
0.25	0.75	Severe	NoAgreement	Weak	NoResults
0	1	Mild	GoverningRules	Strong	GlobalSuccess
0	1	Mild	GoverningRules	Strong	SuccessforUS
0	1	Mild	GoverningRules	Strong	SuccessforOthers
0	1	Mild	GoverningRules	Strong	Failure
0	1	Mild	GoverningRules	Strong	NoResults
0	1	Mild	GoverningRules	Weak	GlobalSuccess
0	1	Mild	GoverningRules	Weak	SuccessforUS
0	1	Mild	GoverningRules	Weak	SuccessforOthers
0	1	Mild	GoverningRules	Weak	Failure
0	1	Mild	GoverningRules	Weak	NoResults
0	1	Mild	Ban	Strong	GlobalSuccess
0	1	Mild	Ban	Strong	SuccessforUS
0	1	Mild	Ban	Strong	SuccessforOthers
0	1	Mild	Ban	Strong	Failure
0	1	Mild	Ban	Strong	NoResults
0	1	Mild	Ban	Weak	GlobalSuccess
0	1	Mild	Ban	Weak	SuccessforUS
0	1	Mild	Ban	Weak	SuccessforOthers
0	1	Mild	Ban	Weak	Failure
0	1	Mild	Ban	Weak	NoResults
0	1	Mild	NoAgreement	Strong	GlobalSuccess
0	1	Mild	NoAgreement	Strong	SuccessforUS
0	1	Mild	NoAgreement	Strong	SuccessforOthers
0	1	Mild	NoAgreement	Strong	Failure
0	1	Mild	NoAgreement	Strong	NoResults
0	1	Mild	NoAgreement	Weak	GlobalSuccess
0	1	Mild	NoAgreement	Weak	SuccessforUS
0	1	Mild	NoAgreement	Weak	SuccessforOthers
0	1	Mild	NoAgreement	Weak	Failure
0	1	Mild	NoAgreement	Weak	NoResults

Table C.7
Conditional Probability Table: USTests

Multilateral	Unilateral	NoTest	Norms
1	0	0	Strong
0	0	1	Ban
0.5	0.5	0	NoNorm

Table C.8
Conditional Probability Table: DeploymentResults

GlobalSuccess	SuccessforUS	SuccessforOthers	Failure	NoResults	GeoEngineeringPotential	TestResults	USDeploys	OthersDeploy
1	0	0	0	0	Likely	GlobalSuccess	Multilateral	Unilateral
1	0	0	0	0	Likely	GlobalSuccess	Multilateral	NoDeployment
0	0.25	0.25	0.5	0	Likely	GlobalSuccess	Unilateral	Unilateral
0	0.42857	0.57143	0	0	Likely	GlobalSuccess	Unilateral	NoDeployment
0	0.57143	0.42857	0	0	Likely	GlobalSuccess	NoDeployment	Unilateral
0	0	0	0	1	Likely	GlobalSuccess	NoDeployment	NoDeployment
1	0	0	0	0	Likely	SuccessforUS	Multilateral	Unilateral
1	0	0	0	0	Likely	SuccessforUS	Multilateral	NoDeployment
0	0.45455	0.04545	0.5	0	Likely	SuccessforUS	Unilateral	Unilateral
0	0.98684	0.01316	0	0	Likely	SuccessforUS	Unilateral	NoDeployment
0	0.57143	0.42857	0	0	Likely	SuccessforUS	NoDeployment	Unilateral
0	0	0	0	1	Likely	SuccessforUS	NoDeployment	NoDeployment
1	0	0	0	0	Likely	SuccessforOthers	Multilateral	Unilateral
1	0	0	0	0	Likely	SuccessforOthers	Multilateral	NoDeployment
0	0.04545	0.45455	0.5	0	Likely	SuccessforOthers	Unilateral	Unilateral
0	0.42857	0.57143	0	0	Likely	SuccessforOthers	Unilateral	NoDeployment
0	0.01316	0.98684	0	0	Likely	SuccessforOthers	NoDeployment	Unilateral
0	0	0	0	1	Likely	SuccessforOthers	NoDeployment	NoDeployment
0.1	0.045	0.045	0.81	0	Likely	Failure	Multilateral	Unilateral
0.1	0.045	0.045	0.81	0	Likely	Failure	Multilateral	NoDeployment
0	0.025	0.025	0.95	0	Likely	Failure	Unilateral	Unilateral

Table C.8—Continued

GlobalSuccess	SuccessforUS	SuccessforOthers	Failure	NoResults	GeoEngineeringPotential	TestResults	USDeploys	OthersDeploy
0	0.04286	0.05714	0.9	0	Likely	Failure	Unilateral	NoDeployment
0	0.05714	0.04286	0.9	0	Likely	Failure	NoDeployment	Unilateral
0	0	0	0	1	Likely	Failure	NoDeployment	NoDeployment
0.3	0.105	0.105	0.49	0	Likely	NoResults	Multilateral	Unilateral
0.3	0.105	0.105	0.49	0	Likely	NoResults	Multilateral	NoDeployment
0	0.075	0.075	0.85	0	Likely	NoResults	Unilateral	Unilateral
0	0.12857	0.17143	0.7	0	Likely	NoResults	Unilateral	NoDeployment
0	0.17143	0.12857	0.7	0	Likely	NoResults	NoDeployment	Unilateral
0	0	0	0	1	Likely	NoResults	NoDeployment	NoDeployment
0.25	0.125	0.125	0.5	0	Unlikely	GlobalSuccess	Multilateral	Unilateral
0.25	0.125	0.125	0.5	0	Unlikely	GlobalSuccess	Multilateral	NoDeployment
0	0.0625	0.0625	0.875	0	Unlikely	GlobalSuccess	Unilateral	Unilateral
0	0.10714	0.14286	0.75	0	Unlikely	GlobalSuccess	Unilateral	NoDeployment
0	0.14286	0.10714	0.75	0	Unlikely	GlobalSuccess	NoDeployment	Unilateral
0	0	0	0	1	Unlikely	GlobalSuccess	NoDeployment	NoDeployment
0.25	0.125	0.125	0.5	0	Unlikely	SuccessforUS	Multilateral	Unilateral
0.25	0.125	0.125	0.5	0	Unlikely	SuccessforUS	Multilateral	NoDeployment
0	0.11364	0.01136	0.875	0	Unlikely	SuccessforUS	Unilateral	Unilateral
0	0.24671	0.00329	0.75	0	Unlikely	SuccessforUS	Unilateral	NoDeployment
0	0.14286	0.10714	0.75	0	Unlikely	SuccessforUS	NoDeployment	Unilateral
0	0	0	0	1	Unlikely	SuccessforUS	NoDeployment	NoDeployment
0.25	0.125	0.125	0.5	0	Unlikely	SuccessforOthers	Multilateral	Unilateral

Table C.8—Continued

GlobalSuccess	SuccessforUS	SuccessforOthers	Failure	NoResults	GeoEngineeringPotential	TestResults	USDeploys	OthersDeploy
0.25	0.125	0.125	0.5	0	Unlikely	SuccessforOthers	Multilateral	NoDeployment
0	0.01136	0.11364	0.875	0	Unlikely	SuccessforOthers	Unilateral	Unilateral
0	0.10714	0.14286	0.75	0	Unlikely	SuccessforOthers	Unilateral	NoDeployment
0	0.00329	0.24671	0.75	0	Unlikely	SuccessforOthers	NoDeployment	Unilateral
0	0	0	0	1	Unlikely	SuccessforOthers	NoDeployment	NoDeployment
0.025	0.0125	0.0125	0.95	0	Unlikely	Failure	Multilateral	Unilateral
0.025	0.0125	0.0125	0.95	0	Unlikely	Failure	Multilateral	NoDeployment
0	0.00625	0.00625	0.9875	0	Unlikely	Failure	Unilateral	Unilateral
0	0.01071	0.01429	0.975	0	Unlikely	Failure	Unilateral	NoDeployment
0	0.01429	0.01071	0.975	0	Unlikely	Failure	NoDeployment	Unilateral
0	0	0	0	1	Unlikely	Failure	NoDeployment	NoDeployment
0.075	0.0375	0.0375	0.85	0	Unlikely	NoResults	Multilateral	Unilateral
0.075	0.0375	0.0375	0.85	0	Unlikely	NoResults	Multilateral	NoDeployment
0	0.01875	0.01875	0.9625	0	Unlikely	NoResults	Unilateral	Unilateral
0	0.03214	0.04286	0.925	0	Unlikely	NoResults	Unilateral	NoDeployment
0	0.04286	0.03214	0.925	0	Unlikely	NoResults	NoDeployment	Unilateral
0	0	0	0	1	Unlikely	NoResults	NoDeployment	NoDeployment

Table C.9
Conditional Probability Table: ClimateImpactsUS

High	Low	EmissionsReductions	ClimateChange	DeploymentResults
0.125	0.875	Strong	Catastrophic	GlobalSuccess
0.125	0.875	Strong	Catastrophic	SuccessforUS
1	0	Strong	Catastrophic	SuccessforOthers
1	0	Strong	Catastrophic	Failure
0.25	0.75	Strong	Catastrophic	NoResults
0.0625	0.9375	Strong	Severe	GlobalSuccess
0.0625	0.9375	Strong	Severe	SuccessforUS
1	0	Strong	Severe	SuccessforOthers
1	0	Strong	Severe	Failure
0.125	0.875	Strong	Severe	NoResults
0	1	Strong	Mild	GlobalSuccess
0	1	Strong	Mild	SuccessforUS
1	0	Strong	Mild	SuccessforOthers
1	0	Strong	Mild	Failure
0	1	Strong	Mild	NoResults
0.5	0.5	Weak	Catastrophic	GlobalSuccess
0.5	0.5	Weak	Catastrophic	SuccessforUS
1	0	Weak	Catastrophic	SuccessforOthers
1	0	Weak	Catastrophic	Failure
1	0	Weak	Catastrophic	NoResults
0.25	0.75	Weak	Severe	GlobalSuccess
0.25	0.75	Weak	Severe	SuccessforUS
1	0	Weak	Severe	SuccessforOthers
1	0	Weak	Severe	Failure
0.5	0.5	Weak	Severe	NoResults
0	1	Weak	Mild	GlobalSuccess
0	1	Weak	Mild	SuccessforUS
1	0	Weak	Mild	SuccessforOthers
1	0	Weak	Mild	Failure
0	1	Weak	Mild	NoResults

Table C.10
Conditional Probability Table: CollateralDamage

High	Low	EmissionsReductions	ClimateChange	DeploymentResults
0.125	0.875	Strong	Catastrophic	GlobalSuccess
1	0	Strong	Catastrophic	SuccessforUS
1	0	Strong	Catastrophic	SuccessforOthers
1	0	Strong	Catastrophic	Failure
0.25	0.75	Strong	Catastrophic	NoResults
0.0625	0.9375	Strong	Severe	GlobalSuccess
1	0	Strong	Severe	SuccessforUS
1	0	Strong	Severe	SuccessforOthers
1	0	Strong	Severe	Failure
0.125	0.875	Strong	Severe	NoResults
0	1	Strong	Mild	GlobalSuccess
1	0	Strong	Mild	SuccessforUS
1	0	Strong	Mild	SuccessforOthers
1	0	Strong	Mild	Failure
0	1	Strong	Mild	NoResults
0.5	0.5	Weak	Catastrophic	GlobalSuccess
1	0	Weak	Catastrophic	SuccessforUS
1	0	Weak	Catastrophic	SuccessforOthers
1	0	Weak	Catastrophic	Failure
1	0	Weak	Catastrophic	NoResults
0.25	0.75	Weak	Severe	GlobalSuccess
1	0	Weak	Severe	SuccessforUS
1	0	Weak	Severe	SuccessforOthers
1	0	Weak	Severe	Failure
0.5	0.5	Weak	Severe	NoResults
0	1	Weak	Mild	GlobalSuccess
1	0	Weak	Mild	SuccessforUS
1	0	Weak	Mild	SuccessforOthers
1	0	Weak	Mild	Failure
0	1	Weak	Mild	NoResults

Table C.11
Conditional Probability Table: Conflict

State	ClimateImpactsUS	ClimateImpactsOthers	CollateralDamage	DeploymentResults
Low	High	High	High	GlobalSuccess
Low	High	High	High	SuccessforUS
Low	High	High	High	SuccessforOthers
High	High	High	High	Failure
Low	High	High	High	NoResults
Low	High	High	Low	GlobalSuccess
Low	High	High	Low	SuccessforUS
Low	High	High	Low	SuccessforOthers
High	High	High	Low	Failure
Low	High	High	Low	NoResults
High	High	Low	High	GlobalSuccess
High	High	Low	High	SuccessforUS
High	High	Low	High	SuccessforOthers
High	High	Low	High	Failure
Low	High	Low	High	NoResults
High	High	Low	Low	GlobalSuccess
High	High	Low	Low	SuccessforUS
High	High	Low	Low	SuccessforOthers
High	High	Low	Low	Failure
Low	High	Low	Low	NoResults
High	Low	High	High	GlobalSuccess
High	Low	High	High	SuccessforUS
High	Low	High	High	SuccessforOthers
High	Low	High	High	Failure
Low	Low	High	High	NoResults
High	Low	High	Low	GlobalSuccess
High	Low	High	Low	SuccessforUS
High	Low	High	Low	SuccessforOthers
High	Low	High	Low	Failure
Low	Low	High	Low	NoResults
Low	Low	Low	High	GlobalSuccess
Low	Low	Low	High	SuccessforUS
Low	Low	Low	High	SuccessforOthers

Table C.11—Continued

State	ClimateImpactsUS	ClimateImpactsOthers	CollateralDamage	DeploymentResults
High	Low	Low	High	Failure
Low	Low	Low	High	NoResults
Low	Low	Low	Low	GlobalSuccess
Low	Low	Low	Low	SuccessforUS
Low	Low	Low	Low	SuccessforOthers
High	Low	Low	Low	Failure
Low	Low	Low	Low	NoResults

Table C.12
Conditional Probability Table: ClimateImpactsOthers

High	Low	EmissionsReductions	ClimateChange	DeploymentResults
0.125	0.875	Strong	Catastrophic	GlobalSuccess
1	0	Strong	Catastrophic	SuccessforUS
0.125	0.875	Strong	Catastrophic	SuccessforOthers
1	0	Strong	Catastrophic	Failure
0.25	0.75	Strong	Catastrophic	NoResults
0.0625	0.9375	Strong	Severe	GlobalSuccess
1	0	Strong	Severe	SuccessforUS
0.0625	0.9375	Strong	Severe	SuccessforOthers
1	0	Strong	Severe	Failure
0.125	0.875	Strong	Severe	NoResults
0	1	Strong	Mild	GlobalSuccess
1	0	Strong	Mild	SuccessforUS
0	1	Strong	Mild	SuccessforOthers
1	0	Strong	Mild	Failure
0	1	Strong	Mild	NoResults
0.5	0.5	Weak	Catastrophic	GlobalSuccess
1	0	Weak	Catastrophic	SuccessforUS
0.5	0.5	Weak	Catastrophic	SuccessforOthers
1	0	Weak	Catastrophic	Failure
1	0	Weak	Catastrophic	NoResults
0.25	0.75	Weak	Severe	GlobalSuccess
1	0	Weak	Severe	SuccessforUS
0.25	0.75	Weak	Severe	SuccessforOthers
1	0	Weak	Severe	Failure
0.5	0.5	Weak	Severe	NoResults
0	1	Weak	Mild	GlobalSuccess
1	0	Weak	Mild	SuccessforUS
0	1	Weak	Mild	SuccessforOthers
1	0	Weak	Mild	Failure
0	1	Weak	Mild	NoResults

Table C.13
Conditional Probability Table: USDeploys

Multilateral	Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
1	0	0	Catastrophic	GoverningRules	Strong	GlobalSuccess
0	1	0	Catastrophic	GoverningRules	Strong	SuccessforUS
0	0	1	Catastrophic	GoverningRules	Strong	SuccessforOthers
0	0	1	Catastrophic	GoverningRules	Strong	Failure
0	1	0	Catastrophic	GoverningRules	Strong	NoResults
1	0	0	Catastrophic	GoverningRules	Weak	GlobalSuccess
0	1	0	Catastrophic	GoverningRules	Weak	SuccessforUS
0	0	1	Catastrophic	GoverningRules	Weak	SuccessforOthers
0	0	1	Catastrophic	GoverningRules	Weak	Failure
0	1	0	Catastrophic	GoverningRules	Weak	NoResults
0	1	0	Catastrophic	Ban	Strong	GlobalSuccess
0	1	0	Catastrophic	Ban	Strong	SuccessforUS
0	0	1	Catastrophic	Ban	Strong	SuccessforOthers
0	0	1	Catastrophic	Ban	Strong	Failure
0	1	0	Catastrophic	Ban	Strong	NoResults
0	1	0	Catastrophic	Ban	Weak	GlobalSuccess
0	1	0	Catastrophic	Ban	Weak	SuccessforUS
0	0	1	Catastrophic	Ban	Weak	SuccessforOthers
0	0	1	Catastrophic	Ban	Weak	Failure
0	1	0	Catastrophic	Ban	Weak	NoResults
0	1	0	Catastrophic	NoAgreement	Strong	GlobalSuccess

Table C.13—Continued

Multilateral	Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
0	1	0	Catastrophic	NoAgreement	Strong	SuccessforUS
0	0	1	Catastrophic	NoAgreement	Strong	SuccessforOthers
0	0	1	Catastrophic	NoAgreement	Strong	Failure
0	1	0	Catastrophic	NoAgreement	Strong	NoResults
0	1	0	Catastrophic	NoAgreement	Weak	GlobalSuccess
0	1	0	Catastrophic	NoAgreement	Weak	SuccessforUS
0	0	1	Catastrophic	NoAgreement	Weak	SuccessforOthers
0	0	1	Catastrophic	NoAgreement	Weak	Failure
0	1	0	Catastrophic	NoAgreement	Weak	NoResults
0.25	0	0.75	Severe	GoverningRules	Strong	GlobalSuccess
0	0	1	Severe	GoverningRules	Strong	SuccessforUS
0	0	1	Severe	GoverningRules	Strong	SuccessforOthers
0	0	1	Severe	GoverningRules	Strong	Failure
0	0	1	Severe	GoverningRules	Strong	NoResults
1	0	0	Severe	GoverningRules	Weak	GlobalSuccess
0	0	1	Severe	GoverningRules	Weak	SuccessforUS
0	0	1	Severe	GoverningRules	Weak	SuccessforOthers
0	0	1	Severe	GoverningRules	Weak	Failure
0	0	1	Severe	GoverningRules	Weak	NoResults
0	0	1	Severe	Ban	Strong	GlobalSuccess
0	0	1	Severe	Ban	Strong	SuccessforUS
0	0	1	Severe	Ban	Strong	SuccessforOthers

Table C.13—Continued

Multilateral	Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
0	0	1	Severe	Ban	Strong	Failure
0	0	1	Severe	Ban	Strong	NoResults
0	0	1	Severe	Ban	Weak	GlobalSuccess
0	0	1	Severe	Ban	Weak	SuccessforUS
0	0	1	Severe	Ban	Weak	SuccessforOthers
0	0	1	Severe	Ban	Weak	Failure
0	0	1	Severe	Ban	Weak	NoResults
0.125	0.25	0.625	Severe	NoAgreement	Strong	GlobalSuccess
0	0.1875	0.8125	Severe	NoAgreement	Strong	SuccessforUS
0	0	1	Severe	NoAgreement	Strong	SuccessforOthers
0	0	1	Severe	NoAgreement	Strong	Failure
0	0.0625	0.9375	Severe	NoAgreement	Strong	NoResults
0.5	0.5	0	Severe	NoAgreement	Weak	GlobalSuccess
0	0.75	0.25	Severe	NoAgreement	Weak	SuccessforUS
0	0	1	Severe	NoAgreement	Weak	SuccessforOthers
0	0	1	Severe	NoAgreement	Weak	Failure
0	0.25	0.75	Severe	NoAgreement	Weak	NoResults
0	0	1	Mild	GoverningRules	Strong	GlobalSuccess
0	0	1	Mild	GoverningRules	Strong	SuccessforUS
0	0	1	Mild	GoverningRules	Strong	SuccessforOthers
0	0	1	Mild	GoverningRules	Strong	Failure
0	0	1	Mild	GoverningRules	Strong	NoResults

Table C.13—Continued

Multilateral	Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
0	0	1	Mild	GoverningRules	Weak	GlobalSuccess
0	0	1	Mild	GoverningRules	Weak	SuccessforUS
0	0	1	Mild	GoverningRules	Weak	SuccessforOthers
0	0	1	Mild	GoverningRules	Weak	Failure
0	0	1	Mild	GoverningRules	Weak	NoResults
0	0	1	Mild	Ban	Strong	GlobalSuccess
0	0	1	Mild	Ban	Strong	SuccessforUS
0	0	1	Mild	Ban	Strong	SuccessforOthers
0	0	1	Mild	Ban	Strong	Failure
0	0	1	Mild	Ban	Strong	NoResults
0	0	1	Mild	Ban	Weak	GlobalSuccess
0	0	1	Mild	Ban	Weak	SuccessforUS
0	0	1	Mild	Ban	Weak	SuccessforOthers
0	0	1	Mild	Ban	Weak	Failure
0	0	1	Mild	Ban	Weak	NoResults
0	0	1	Mild	NoAgreement	Strong	GlobalSuccess
0	0	1	Mild	NoAgreement	Strong	SuccessforUS
0	0	1	Mild	NoAgreement	Strong	SuccessforOthers
0	0	1	Mild	NoAgreement	Strong	Failure
0	0	1	Mild	NoAgreement	Strong	NoResults
0	0	1	Mild	NoAgreement	Weak	GlobalSuccess
0	0	1	Mild	NoAgreement	Weak	SuccessforUS

Table C.13—Continued

Multilateral	Unilateral	NoDeployment	ClimateChange	DeploymentAgreement	EmissionsReductions	TestResults
0	0	1	Mild	NoAgreement	Weak	SuccessforOthers
0	0	1	Mild	NoAgreement	Weak	Failure
0	0	1	Mild	NoAgreement	Weak	NoResults

References

- Airports Council International, “Statistics: Top 10 World Airports,” information brief, 2009. As of September 11, 2010:
http://www.airports.org/aci/aci/file/Position%20Briefs/position%20brief_TOP10_2009.pdf
- Bala, G., P. B. Duffy, and K. E. Taylor, “Impact of Geoengineering Schemes on the Global Hydrological Cycle,” *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 105, No. 22, June 3, 2008, pp. 7664–7669.
- Bankes, Steve, “Exploratory Modeling for Policy Analysis,” *Operations Research*, Vol. 41, No. 3, May–June 1993, pp. 435–449.
- Barrett, Scott, “The Incredible Economics of Geoengineering,” *Environmental and Resource Economics*, Vol. 39, No. 1, 2008, pp. 45–54.
- Bernstein, Lenny, Rajendra K. Pachauri, and Andy Reisinger, eds., *Climate Change 2007: Synthesis Report*, Geneva: Intergovernmental Panel on Climate Change, 2008. As of September 23, 2010:
http://www.ipcc.ch/publications_and_data/ar4/syr/en/contents.html
- Bickel, J. Eric, and Lee Lane, *An Analysis of Climate Engineering As a Response to Climate Change*, Frederiksberg, Denmark: Copenhagen Consensus Center, September 2009. As of September 11, 2010:
http://fixthecclimate.com/fileadmin/templates/page/scripts/downloadpdf.php?file=/uploads/tx_templavoila/AP_Climate_Engineering_Bickel_Lane_v.5.0.pdf
- Blackstock, Jason J., David S. Battisti, Ken Caldeira, Douglas M. Eardley, Jonathan I. Katz, David W. Keith, Aristides A. N. Patrinos, Daniel P. Schrag, Robert H. Socolow, and Steven E. Koonin, *Climate Engineering Responses to Climate Emergencies*, Santa Barbara, Calif.: Novim, July 28, 2009. As of September 11, 2010:
http://www.novim.org/attachments/031_037_Novim%20Report%20Final%2007.28.09.pdf
- Blackstock, Jason J., and Jane C. S. Long, “The Politics of Geoengineering,” *Science*, Vol. 327, No. 5965, January 29, 2010, p. 527.
- Boyd, Philip W., “Ranking Geo-Engineering Schemes,” *Nature Geoscience*, Vol. 1, 2008, pp. 722–724.
- Bryant, Benjamin P., and Robert J. Lempert, “Thinking Inside the Box: A Participatory, Computer-Assisted Approach to Scenario Discovery,” *Technology Forecasting and Social Change*, Vol. 77, No. 1, January 2010, pp. 34–49.
- Cain, Jeremy, *Planning Improvements in Natural Resources Management: Guidelines for Using Bayesian Networks to Support the Planning and Management of Development Programmes in the Water Sector and Beyond*, Wallingford, Oxon: Centre for Ecology and Hydrology, 2001.
- Caldeira, Ken, and Lowell Wood, “Global and Arctic Climate Engineering: Numerical Model Studies,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical, and Engineering Sciences*, Vol. 366, No. 1882, November 2008, pp. 4039–4056.
- Chayes, Abram, and Antonia Handler Chayes, *The New Sovereignty: Compliance with International Regulatory Agreements*, Cambridge, Mass.: Harvard University Press, 1998.
- Chu, Steven, secretary, U.S. Department of Energy, presentation to employees of SLAC National Accelerator Laboratory, Stanford University, June 27, 2009. As of September 23, 2010:
<http://today.slac.stanford.edu/feature/2009/chu-20090626-video.asp>

Cicerone, Ralph J., "Geoengineering: Encouraging Research and Overseeing Implementation," *Climatic Change*, Vol. 77, No. 3–4, 2006, pp. 221–226.

Cicerone, Ralph J., Scott Elliott, and Richard P. Turco, "Global Environmental Engineering," *Nature*, Vol. 356, No. 6369, April 9, 1992, p. 472.

Clark, Peter U., Andrew J. Weaver, Konrad Steffen, E. R. Cook, and Edward J. Brook, *Abrupt Climate Change: Final Report*, Washington, D.C.: U.S. Climate Change Science Program, synthesis and assessment product 3.4, December 2008. As of September 11, 2010:
<http://purl.access.gpo.gov/GPO/LPS110700>

Convention on Biological Diversity, home page, updated March 4, 2011. As of March 7, 2011:
<http://www.cbd.int/>

Convention on the Prohibition of Military or Any Hostile Use of Environmental Modification Techniques, December 10, 1976. As of March 7, 2011:
<http://www.icrc.org/ihl.nsf/FULL/460>

Dessai, Suraje, Mike Hulme, Robert Lempert, and Roger Pielke Jr., "Do We Need Better Predictions to Adapt to a Changing Climate?" *EOS*, Vol. 90, No. 13, 2009, pp. 111–112.

Dewar, James A., *Assumption-Based Planning: A Tool for Reducing Avoidable Surprises*, Cambridge: Cambridge University Press, 2002.

Doney, Scott C., Victoria J. Fabry, Richard A. Feely, and Joan A. Kleypas, "Ocean Acidification: The Other CO₂ Problem," *Annual Review of Marine Science*, Vol. 1, January 2009, pp. 169–192.

Fleming, James R., "The Climate Engineers," *Wilson Quarterly*, Spring 2007.

Goes, Marlos, Klaus Keller, and Nancy Tuana, "The Economics (or Lack Thereof) of Aerosol Geoengineering," submitted April 2010. As of September 23, 2010:
http://www.geosc.psu.edu/~kkeller/Goes_et_al_geoengineering_cc_2009_submitted.pdf

Groves, David G., and Robert J. Lempert, "A New Analytic Method for Finding Policy-Relevant Scenarios," *Global Environmental Change*, Vol. 17, No. 1, 2007, pp. 78–85. As of March 7, 2011:
<http://www.rand.org/pubs/reprints/RP1244.html>

Haas, Peter M., "Introduction: Epistemic Communities and International Policy Coordination," *International Organization*, Vol. 46, No. 1, Winter 1992, pp. 1–35.

Hansen, James, Makiko Sato, Pushker Kharecha, David Beerling, Valerie Masson-Delmotte, Mark Pagani, Maureen Raymo, Dana L. Royer, and James C. Zachos, "Target Atmospheric CO₂: Where Should Humanity Aim?" *Open Atmospheric Science Journal*, Vol. 2, 2008, pp. 217–231.

House, Kurt Zenz, Christopher H. House, Daniel P. Schrag, and Michael J. Aziz, "Electrochemical Acceleration of Chemical Weathering as an Energetically Feasible Approach to Mitigating Anthropogenic Climate Change," *Environmental Science and Technology*, Vol. 41, No. 24, 2007, pp. 8464–8470.

International Maritime Organization, Convention on the Prevention of Marine Pollution by Dumping of Wastes and Other Matter, 1972 and 1996 Protocol Thereto, "Statement of Concern Regarding Iron Fertilization of the Oceans to Sequester CO₂," July 13, 2007. As of September 11, 2010:
http://www.who.edu/cms/files/London_Convention_statement_24743_29324.pdf

———, "On the Regulation of Ocean Fertilization," resolution LC-LP.1(2008), October 31, 2008. As of September 11, 2010:
<http://ioc3.unesco.org/oanet/OAdocs/LC-LPResOF.pdf>

Izrael, Yu. A., V. M. Zakharov, N. N. Petrov, A. G. Ryaboshapko, V. N. Ivanov, A. V. Savchenko, Yu. V. Andreev, V. G. Eran'kov, Yu. A. Puzov, B. G. Danilyan, V. P. Kulyapin, and V. A. Gulevskii, "Field Studies of a Geo-Engineering Method of Maintaining a Modern Climate with Aerosol Particles," *Russian Meteorology and Hydrology*, Vol. 34, No. 10, 2009, pp. 635–638.

Jackman, Simon, "Bayesian Analysis for Political Research," *Annual Review of Political Science*, Vol. 7, June 2004, pp. 483–505.

Keith, David W., "Geoengineering the Climate: History and Prospect," *Annual Review of Energy and the Environment*, Vol. 25, November 2000, pp. 245–284.

- Keith, David W., Edward Parson, and M. Granger Morgan, "Research on Global Sun Block Needed Now," *Nature*, Vol. 463, No. 7280, January 28, 2010, pp. 426–428.
- Kintisch, Eli, "Asilomar 2' Takes Small Steps Toward Rules for Geoengineering," *Science*, Vol. 328, No. 5974, April 2, 2010, pp. 22–23.
- Knutti, Reto, "Should We Believe Model Predictions of Future Climate Change?" *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 366, No. 1885, December 2008, pp. 4647–4664.
- Lempert, Robert J., "Can Scenarios Help Policymakers Be Both Bold and Careful?" in Francis Fukuyama, ed., *Blindside: How to Anticipate Forcing Events and Wild Cards in Global Politics*, Washington, D.C.: Brookings Institution Press, 2007, pp. 109–119.
- , "Setting Appropriate Goals: A Priority, Long-Term Climate-Policy Decision," in Robert J. Lempert, Steven W. Popper, Endy Y. Min, and James A. Dewar, eds., *Shaping Tomorrow Today: Near-Term Steps Towards Long-Term Goals*, Santa Monica, Calif.: RAND Corporation, CR-267-RPC, 2009, pp. 63–78. As of March 7, 2011:
http://www.rand.org/pubs/conf_proceedings/CF267.html
- Lempert, Robert J., and Myles T. Collins, "Managing the Risk of Uncertain Threshold Responses: Comparison of Robust, Optimum, and Precautionary Approaches," *Risk Analysis*, Vol. 27, No. 4, 2007, pp. 1009–1026.
- Lempert, Robert J., David G. Groves, Steven W. Popper, and Steven C. Bankes, "A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios," *Management Science*, Vol. 52, No. 4, April 2006, pp. 514–528.
- Lempert, Robert, Nebojsa Nakicenovic, Daniel Sarewitz, and Michael Schlesinger, "Characterizing Climate-Change Uncertainties for Decision-Makers: An Editorial Essay," *Climatic Change*, Vol. 65, No. 1–2, 2004, pp. 1–9.
- Lempert, Robert J., Steven W. Popper, and Steven C. Bankes, "Confronting Surprise," *Social Science Computer Review*, Vol. 20, No. 4, Winter 2002, pp. 420–440.
- , *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*, Santa Monica, Calif.: RAND Corporation, MR-1626-RPC, 2003. As of March 7, 2011:
http://www.rand.org/pubs/monograph_reports/MR1626.html
- Lempert, Robert J., Steven W. Popper, Susan A. Resetar, and Stuart L. Hart, *Capital Cycles and the Timing of Climate Change Policy*, Arlington, Va.: Pew Center on Global Climate Change, October 2002. As of September 11, 2010:
<http://www.pewclimate.org/publications/report/capital-cycles-and-timing-climate-change-policy>
- Lenton, Timothy M., Hermann Held, Elmar Kriegler, Jim W. Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber, "Tipping Elements in the Earth's Climate System," *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 105, No. 6, February 12, 2008, pp. 1786–1793.
- Lin, Albert C., "Geoengineering Governance," *Issues in Legal Scholarship*, Vol. 8, No. 3, 2009, Art. 2. As of September 11, 2010:
<http://www.bepress.com/ils/vol8/iss3/art2/>
- Marcot, Bruce G., Richard S. Holthausen, Martin G. Raphael, Mary M. Rowland, and Michael J. Wisdom, "Using Bayesian Belief Networks to Evaluate Fish and Wildlife Population Viability Under Land Management Alternatives from an Environmental Impact Statement," *Forest Ecology and Management*, Vol. 153, No. 1–3, October 1, 2001, pp. 29–42.
- McCann, Robert K., Bruce G. Marcot, and Rick Ellis, "Bayesian Belief Networks: Applications in Ecology and Natural Resource Management," *Canadian Journal of Forest Research*, Vol. 36, No. 12, 2006, pp. 3053–3062.

Morgan, Granger, H. Dowlatbadi, M. Henrion, D. Keith, R. Lempert, S. McBrid, M. Small, and T. Wilbanks, eds., *Best Practice Approaches for Characterizing, Communicating, and Incorporating Scientific Uncertainty in Decisionmaking*, Washington, D.C.: National Oceanic and Atmospheric Administration, final report, synthesis and assessment product 5.2, January 2009. As of September 11, 2010: <http://www.climate-science.gov/Library/sap/sap5-2/final-report/default.htm>

National Research Council, Committee on Abrupt Climate Change, *Abrupt Climate Change: Inevitable Surprises*, Washington, D.C.: National Academy Press, 2002.

———, Committee on Stabilization Targets for Atmospheric Greenhouse Gas Concentrations, *Climate Stabilization Targets: Emissions, Concentrations, and Impacts over Decades to Millennia*, Washington, D.C.: National Academies Press, 2010a.

———, Panel on Informing Effective Decisions and Actions Related to Climate Change, Board on Atmospheric Sciences and Climate, *Informing an Effective Response to Climate Change: America's Climate Choices*, Washington, D.C.: National Academies Press, 2010b.

NRC—See National Research Council.

Patashnik, Eric, “After the Public Interest Prevails: The Political Sustainability of Policy Reform,” *Governance*, Vol. 16, No. 2, April 2003, pp. 203–234.

Polborn, Sarah, and Felix Tintelnot, “How Geoengineering May Encourage Carbon Dioxide Abatement,” State College, Pa.: Center for Research on International Financial and Energy Security, May 1, 2009. As of September 11, 2010: <http://crifes.psu.edu/papers/polbornpaper.pdf>

RAND Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition, “Pardee Center,” last modified November 23, 2010. As of March 7, 2011: http://www.rand.org/international_programs/pardee.html

Rayfuse, Rosemary, Mark G. Lawrence, and Kristina M. Gjerde, “Ocean Fertilisation and Climate Change: The Need to Regulate Emerging High Seas Uses,” *International Journal of Marine and Coastal Law*, Vol. 23, No. 2, 2008, pp. 297–326.

Robock, Alan, Martin Bunzi, Ben Kravitz, and Georgiy L. Stenchikov, “A Test for Geoengineering?” *Science*, Vol. 327, No. 5965, January 29, 2010, pp. 530–531.

Rockström, Johan, Will Steffen, Kevin Noone, Åsa Persson, F. Stuart Chapin III, Eric F. Lambin, Timothy M. Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, Björn Nykvist, Cynthia A. de Wit, Terry Hughes, Sander van der Leeuw, Henning Rodhe, Sverker Sörlin, Peter K. Snyder, Robert Costanza, Uno Svedin, Malin Falkenmark, Louise Karlberg, Robert W. Corell, Victoria J. Fabry, James Hansen, Brian Walker, Diana Liverman, Katherine Richardson, Paul Crutzen, and Jonathan A. Foley, “A Safe Operating Space for Humanity,” *Nature*, Vol. 461, No. 7263, September 24, 2009, pp. 472–476.

Royal Society, *Geoengineering the Climate: Science, Governance and Uncertainty*, London, September 1, 2009. As of September 11, 2010: <http://royalsociety.org/geoengineering-the-climate/>

Saltelli, A., K. Chan, and E. Marian Scott, eds., *Sensitivity Analysis*, Chichester: Wiley, 2000.

Savage, Leonard J., *The Foundations of Statistics*, New York: Wiley, 1954.

Schelling, Thomas C., “The Economic Diplomacy of Geoengineering,” *Climatic Change*, Vol. 33, No. 3, 1996, pp. 303–307.

Schneider, Stephen H., “Geoengineering: Could—or Should—We Do It?” *Climatic Change*, Vol. 33, No. 3, 1996, pp. 291–302.

Slaughter, Anne-Marie, *A New World Order*, Princeton, N.J.: Princeton University Press, 2004.

Solomon, Susan, Gian-Kasper Plattner, Reto Knutti, and Pierre Friedlingstein, “Irreversible Climate Change Due to Carbon Dioxide Emissions,” *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 106, No. 6, February 10, 2009, pp. 1704–1709.

Treaty Between the United States of America and the Union of Soviet Socialist Republics on the Limitation of Anti-Ballistic Missile Systems, May 26, 1972.

United Nations Environment Programme, Ozone Secretariat, *The Montreal Protocol on Substances That Deplete the Ozone Layer as Either Adjusted and/or Amended in London 1990, Copenhagen 1992, Vienna 1995, Montreal 1997, Beijing, 1999, Nairobi, 2000*, 2000. As of March 7, 2011:
<http://www.unep.org/ozone/pdfs/montreal-protocol2000.pdf>

Victor, David G., “On the Regulation of Geoengineering,” *Oxford Review of Economic Policy*, Vol. 24, No. 2, 2008, pp. 322–336.

Victor, David G., M. Granger Morgan, Jay Apt, John Steinbruner, and Katharine Ricke, “The Geoengineering Option,” *Foreign Affairs*, Vol. 88, No. 2, 2009, pp. 64–76.

Virgoe, John, “International Governance of a Possible Geoengineering Intervention to Combat Climate Change,” *Climatic Change*, Vol. 95, No. 1–2, 2008, pp. 103–119.

Weitzman, Martin L., “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics*, Vol. 91, No. 1, February 2009, pp. 1–19.