ESSAYS ON UNCERTAINTY ANALYSIS IN ENERGY MODELING: CAPACITY PLANNING, R&D PORTFOLIO MANAGEMENT, AND FAT-TAILED UNCERTAINTY

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MANAGEMENT
SCIENCE AND ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

John Erik Bistline
November 2013

This dissertation is online at: http://purl.stanford.edu/sc063nw4945

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

John Weyant, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Gerd Infanger

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

James Sweeney

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

The characterization and analysis of uncertainty are central components of decision-making, especially in the energy sector; however, there is currently a gap in the energy modeling community between the recognition of uncertainty's importance and its incorporation in large-scale models. This dissertation explores how the explicit inclusion of uncertainty through sequential decision-making approaches like stochastic programming can provide insights to energy planners in different domains.

The dissertation first investigates the dynamics of capacity planning and dispatch in the electric power sector under technological, economic, and policy-related uncertainties. Metrics like the expected value of perfect information and the value of the stochastic solution quantify the benefits of reducing uncertainty and of incorporating uncertainty explicitly in modeling efforts. Model results highlight risks associated with shale gas and climate policy, offer policy guidance in these areas, and indicate that planners are likely underestimating the impacts of uncertainty. Hedging and strategic delay are explained in terms of the optionality of energy investments, leading to insights about uncertainty, learning, and irreversibility.

A second application presents a framework for allocating investments across a portfolio of energy technology research and development (R&D) programs, which incorporates uncertainties in the effectiveness of investments and in diffusion markets. This work analyzes how R&D valuations vary in different decision-making settings and shows how wait-and-see valuation approaches, by not explicitly accounting for exogenous market uncertainties, may undervalue the hedging potential of technologies. The results indicate that R&D is more valuable in suboptimal planning and policy environments.

The final section discusses policy and modeling questions about low-probability, high-impact risks in climate change economics. This analysis examines the impacts of fat-tailed uncertainty about the climate sensitivity parameter on near-term abatement using a sequential decision-making framework. The results demonstrate how policy prescriptions from integrated assessment models are highly sensitive to the specifications of uncertainty, learning, and damages. Fat tails alone do not merit stringent mitigation immediately, which also requires strongly convex damages and slow learning. The analysis illustrates the potential value of midcourse corrections on reducing consumption risks imposed by uncertain damages and focuses attention on the dynamics of learning.

Acknowledgments

"Human experience, which is constantly contradicting theory, is the great test of truth. A system, built upon the discoveries of a great many minds, is always of more strength, than what is produced by the mere workings of any one mind, which, of itself, can do little. There is not so poor a book in the world that would not be a prodigious effort were it wrought out entirely by a single mind, without the aid of prior investigators."

—Samuel Johnson (Boswell's Life of Samuel Johnson)

This dissertation embodies the wisdom, influence, and encouragement of many people. I would like to thank all of them for making this work possible and enjoyable, though only a few will be mentioned here by name.

I am indebted to my advisers for guiding my development as a scholar. It has been a privilege to have John Weyant as my primary adviser and friend. As a researcher, there is no greater gift than the time, space, and support to pursue interesting research questions. I am grateful for the opportunity and honor to work with an adviser who can make new things familiar and familiar things new. John taught me how to think through problems thoroughly and, more important, how to frame worthwhile problems in the first place. Gerd Infanger has been an instrumental member of my reading committee and has provided a bridge to state-of-the-art stochastic programming research. I thank him for offering valuable suggestions about my work with boundless enthusiasm. His contributions considerably increased the level of sophistication of the analysis. Additionally, I am grateful for the guidance of Jim Sweeney who has been a valuable adviser and reading committee member. It has been a pleasure to learn his elegant approach to problem solving. He has helped me to better scrutinize the bark on trees while never forgetting about the surrounding forest.

I thank Yinyu Ye and Chris Edwards for serving on my oral defense committee. As my master's adviser in the Mechanical Engineering Department, Chris helped me to think like a thermodynamicist and profoundly shaped the way I view energy systems.

There are many other researchers (at Stanford and elsewhere) to whom I owe a substantial intellectual debt for their wise counsel and kindness, including Erin Baker, Sally Benson, Adam Brandt, Sig Hecker, Max Henrion, Hill Huntington, Jon Koomey, Granger Morgan, Vic Niemeyer, Elisabeth Paté-Cornell, Varun Rai, Stefan Reichelstein, Margaret Taylor, Tom Wilson, Detlof von Winterfeldt, and many others. I had the good fortune of interning at the Electric Power Research Institute during my time at Stanford and of working with Geoff Blanford whose insight, generosity, and encouragement influenced my approach to modeling. Additionally, I thank the professors and students in the Mechanical Engineering and Engineering and Public Policy Departments at Carnegie Mellon for sparking my interest in research, helping me build a solid skill set, and inspiring me to pursue interstitial research areas related to energy and the environment. My interests in uncertainty analysis and decision-making in the electric power sector began during my time there and ultimately paved the path toward my dissertation topic.

My dissertation research would not have been possible without the financial support of the William K. Bowes, Jr. Stanford Graduate Fellowship. Portions of this research were supported by the Office of Science (BER), US Department of Energy.

My colleagues played important roles throughout my time at Stanford. I would like to extend my sincere gratitude to Steve Comello, Benjamin Leibowicz, Ansu Sahoo, and Delavane Turner for their valuable feedback on chapters of this dissertation. I am also grateful for the friendship, conversation, and research advice of colleagues, including Mark Budolfson, Danielle Davidian, Blake Francis, Ken Gillingham, Michael Gucwa, Doug Hannah, Sébastien Houde, Ahren Lacy, Pete Larsen, Juan Carlos López, Sankaran Ramakrishnan, Chris Rinaldi, Maria Roumpani, Witi Sachchamarga, Adam Simpson, and Jordan Wilkerson.

I thank my friends at Stanford and elsewhere for assisting me along every step of the way and for shaping my outlook on life, the universe, and everything. In particular, I want to thank David, Elah, Erica, Giancarlo, Graham, Josh, and Kristin for their friendship, conversation, and inspiration.

I am indescribably fortunate to have such an amazingly supportive family, without which I would not be where I am today. Words cannot express my gratitude for my generous and loving parents, David and Norma, who have given me every opportunity to succeed and have provided encouragement in all of my pursuits. I also want to thank my brother, Chris, for his feedback on my dissertation and for his friendship and good humor. I thank my grandparents as well for enabling and encouraging me to follow my passions from an early age.

I owe my most sincere gratitude to my ineffably wonderful fiancée, Ellen. She has been a companion in my joys and a comfort in my sorrows throughout my time at Stanford. Ellen's feedback has improved every facet of my work, and her daily encouragement has buoyed my spirits when I have felt discouraged. Her love, support, and patience have been integral to this dissertation and beyond.

JEB

Stanford, California
November 2013

"For the life of a soul does not consist in the contemplation of one consistent world but in the painful task of unifying (to a greater or less extent) jarring and incompatible ones, and passing, when possible, from two or more discordant viewpoints to a higher which shall somehow include and transmute them."

—T. S. Eliot (Knowledge and Experience in the Philosophy of F. H. Bradley)

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation: Uncertainty in Energy Modeling

Utilities, generators, and other energy planners face many pervasive sources of uncertainty when making near-term decisions. Since energy-sector assets like power plants, transmission facilities, and oil refineries are long-lived and largely irreversible investments, insufficiently characterizing or accounting for uncertainties like fuel prices or environmental policy can impose economic and environmental burdens on a range of stakeholders. Due to the long-lasting effects of energy decisions and also to the complicated economic, technological, and policy-related systems in which they are embedded, there are different manifestations of uncertainty and risk management for stakeholders at all levels, from individual firms to nations.

Given the centrality and complexity of uncertainty in energy and environmental management, there is a need for decision support tools that can provide a greater sense of clarity and that can reduce exposure to significant downside losses while preserving options for upside gains from volatility. However, in the energy modeling community, there is currently a gap between the recognition of uncertainty's importance and its actual incorporation in large-scale models. Many energy models currently use deterministic frameworks, and when uncertainty analysis is performed, it often involves simple methods like sensitivity or scenario analysis (Kann and Weyant, 2000). These methods use different assumptions for uncertain parameters to test the robustness of conclusions but do not typically incorporate probabilistic information.

Another commonly used approach is uncertainty propagation (e.g., Monte Carlo analysis), which involves sampling from distributions for uncertain parameters, propagating them through a deterministic model, and creating output distributions. Since a propagation framework suggests a different optimal strategy for each state of the world, this approach leaves decision-makers in a quandary about how to choose among alternatives before uncertainty is resolved and how to translate results of modeling efforts into actionable near-term insights.[1] Additionally, since the costs associated with uncertainty remain unknown, propagation approaches cannot quantify the relative importance of uncertainties.

In contrast to sensitivity analysis and uncertainty propagation approaches, sequential decision-making frameworks like stochastic programming incorporate uncertainty explicitly and identify hedging strategies that balance the risks of premature action with those of delay. These models determine "optimal" policies in multiple stages based on updated information and offer a more robust treatment of uncertainty than propagation models, which select policies once and do not incorporate learning thereafter.[2] Although early research drew attention to the importance of uncertainty analysis and incorporated sequential decision-making in simple energy models (Morgan and Henrion, 1990; Manne and Richels, 1993; Nordhaus, 1994; Birge and Rosa, 1996), there was subsequently a noticeable dearth of research that applied such approaches to large-scale energy models. As described in the literature review in Chapter 2, the limited studies that have used stochastic programming do not take full advantage of the framework and often do not incorporate recent advances in the field, thus overlooking valuable opportunities to inform challenging decision problems. These shortcomings make the development and application of sequential

[1] For many decision problems, optimal strategies under different scenarios suggest vastly different actions in the immediate future. For instance, Blanford (2013) shows how emissions trajectories in standard deterministic models are highly sensitive to assumptions about negative emissions technologies when performing cost-effectiveness analysis with overshoot to reach stabilization targets. However, the future availability and cost of carbon dioxide removal technologies are far from certain.

[2] Mathematical programming models recommend optimal strategies given a range of assumptions about parameters, probabilities, system constraints, and preferences. As this dissertation emphasizes, the validity of a model recommendation depends strongly on such assumptions.

decision-making frameworks in energy modeling an area of active research.

This dissertation examines how the explicit inclusion of uncertainty through sequential decision-making can provide insights to energy planners in the domains of electric sector capacity planning, research and development (R&D) portfolio management, and climate policy.[3] It aims to identify novel conclusions that cannot be captured in (or would not be apparent using) standard deterministic approaches. In doing so, the tools developed and applied here provide better guidance to answering questions surrounding the characterization, analysis, and communication of uncertainty for energy and environmental systems.

Recent decades have provided improvements in the performance and costs of computational tools, which have facilitated the development of larger and more complex models. The concurrent trend in increased data availability also augments modeling capabilities for treating uncertainty explicitly. Despite these advances, decision-makers and modelers in the energy landscape have been slow to adopt developments in operations research and management science (e.g., algorithms for solving stochastic programs), even though such techniques have been applied lucratively in other fields like finance, operations management, agriculture, and telecommunications (Wallace and Ziemba, 2005). This dissertation demonstrates the many potential benefits of bridging state-of-the-art operations research techniques with energy models to provide prescriptive decision support and to identify blind spots in planning. These modeling and computational advances can be leveraged to take into explicit consideration a wider range of potential futures and to hedge against negative outcomes.

Another contribution of this work is to demonstrate the importance of uncertainty quantification. The applications in this dissertation illustrate how model results are sensitive to input distributions across a range of domains. Since many existing studies rely on *ad-hoc* distributions over parameters of interest, these results suggests that uncertainty quantification should be given increased research attention commensurate with its importance in determining model results and should be a co-equal partner with model building.[4] The dissertation characterizes uncertainty using a range of

[3]In addition to these three applications, the techniques and metrics employed in this dissertation can be used in a wide range of decision contexts under uncertainty in diverse fields.

[4]As literature reviews in subsequent chapters indicate, many studies invoke the Laplace criterion

approaches like statistical analyses, expert elicitations, and econometric modeling. In particular, expert elicitations are discussed at length in Appendix B and are used throughout the dissertation due to their importance in quantifying uncertainty about future cost and performance characteristics of energy technologies.

1.2 Analytical Framework

A central theme of this research is that the explicit inclusion of uncertainty through sequential decision-making can offer better performance and provide important insights that would not be apparent using deterministic models or non-probabilistic uncertainty analysis. Stochastic programming is the sequential decision-making approach used throughout this dissertation to formulate strategies that incorporate robustness, regret, and resilience. As described in Chapter 2, stochastic programming has many attractive features that make it a well-suited mathematical framework for bridging near-term decisions with their long-term implications in the context of high-dimensional energy modeling under multifaceted uncertainty.[5]

Stochastic programming finds hedging strategies for problems with uncertain data, which adapt to changing conditions and updated information. Like most sequential decision-making approaches, it exchanges some degree of optimal performance for reduced exposure to broken assumptions. Although the strategies suggested by this approach are unlikely to be ideally suited for every state of the world, these strategies are optimal *ex ante* and provide resilience and adaptability under a variety of future scenarios. When uncertainty is prevalent and the long-term consequences of near-term decisions are imperfectly understood, hedging strategies allow decision-makers to shape available options, to cope with the unknown, to learn from errors, and to exploit new information as conditions change.

---

and assume equal probabilities for all states of the world, which is the simplest non-informative prior. In the absence of probabilistic information, the Principle of Insufficient Reason suggests that decision-makers will act as if probabilities of different outcomes are equally likely, which is descriptively true in many settings (Luce and Raiffa, 1957; Levi, 1974). However, the applications in this dissertation provide evidence that such assumptions many be incorrect and may provide misleading guidance to decision-makers.

[5] Here, the near-term future refers generally to the decision periods before uncertainty is resolved.

Stochastic programs embody many of the principles behind resilient and adaptive strategies for management under uncertainty (Morgan et al., 2009). Resilient strategies are ones that consider a diverse array of possible futures and that perform reasonably well across this range of plausible realizations. Adaptive strategies are designed to facilitate less costly adjustments once more information is available about the future through time, experience, and research. Adaptive strategies relate to the concept of flexibility, which is the degree to which a strategy can be adapted in the future as more information is known about random variables. The goal of flexible strategies is to keep the expected costs associated with misforecasts about the future and surprises as low as possible.[6] These strategies are orthogonal to the conventional notion of determining a single optimal strategy for all states of the world. Although stochastic programming does not make resilience or adaptability explicit objectives for optimization, aspects of these concepts are implicitly accounted for in the formulation and resulting strategies of stochastic programs.

Stochastic programming provides a convenient mathematical framework for defining and quantifying answers to two important, overarching questions for decision-makers and modelers:

- How much should decision-makers be willing to pay for information about uncertain quantities?
- What is the value of incorporating uncertainty explicitly in the decision-making process instead of using a deterministic approximation?

These questions can be answered by using two related metrics for evaluating the importance of uncertainties—namely, the expected value of perfect information and the value of the stochastic solution. As discussed in Chapter 3.3.3, the expected value of perfect information represents the expected change in the objective function value if perfectly accurate forecasts are available and places an upper bound on a decision-maker's willingness to pay for information-gathering activities. The value of

[6] Due to the use of a subjectivist Bayesian interpretation of probability in this research, surprises refer to gaps between "perceived reality and one's expectations" (Morgan et al., 2009), suggesting that a surprise is a property of an observer or decision-maker.

the stochastic solution quantifies the benefit of explicitly accounting for uncertainty through range of scenarios and has important implications in the process of model construction by identifying the most important uncertainties for explicit inclusion. This metric also helps to prioritize uncertainty characterization efforts and indicates whether extensive and resource-intensive quantification is necessary. In summary, the expected value of perfect information is the cost of being uncertain, wheres the value of the stochastic solution can be viewed as the additional cost of pretending that uncertainty does not exist.

1.3 Research Questions

Although risk cannot be eliminated, it can be proactively managed. Manne and Richels coined the term "greenhouse insurance" to describe the problem of developing strategies to cope with climate change under pervasive uncertainty (Manne and Richels, 1993). Three general forms of greenhouse insurance are: 1. Reducing emissions; 2. Investing in R&D for technologies to reduce emissions at lower costs (i.e., active technological improvement); 3. Performing scientific research to reduce uncertainties about climate change and its impacts. This taxonomy of strategies stresses how uncertainties in climate policy interact with those in technology policy.

Given the interconnectedness of climate change mitigation and technological development, the three linked projects in subsequent chapters explore different facets of greenhouse insurance and related issues in energy modeling. These applications of sequential decision-making make recommendations about how uncertainty should be quantified and analyzed in the domains of electric sector capacity planning, energy technology R&D portfolio management, and climate policy under fat-tailed uncertainty. These research efforts reinforce the need for more judicious consideration of uncertainty in modeling efforts, decision-making, and resource allocation.

1.3.1 Electric Sector Capacity Planning under Uncertainty

Although planning in the United States (US) electric power sector was relatively predictable during the industry's first century, utilities and policy makers must now grapple with many simultaneous challenges: increasingly stringent compliance with an array of environmental policies (e.g., greenhouse gas emissions, regional haze, ozone, and hazardous air pollutants), an aging fleet of generators, sudden changes to the economics of fossil resources due to shale gas, an increased policy emphasis on demand-side management resources, modernization of the electricity grid, higher prices for construction materials (e.g., concrete and steel), and an uncertain economy.[7]

Electric utilities are one of the largest industries in the US, holding assets of over $600 billion and having annual sales above $260 billion (Munson, 2005). Assets in the power sector like power plants are designed to last many decades.[8] Given the likelihood that greenhouse gas emissions will be regulated during the lifetime of new units, the environment in which generators come online and operate may be very different from the one in which they are planned. Therefore, investment decisions should be made with a full understanding of the risks and tradeoffs associated with each alternative and should be based on the best-available information. Otherwise, the industry could be locked into investments that expose a wide range of stakeholders to increased risk for years to come.

The objective of the research in Chapters 3 through 5 is to investigate how technological, economic, and policy-related uncertainties may impact the deployment of supply-side technologies in the US electric power sector. In particular, this research looks at how uncertainties in the stringency of climate policy, natural gas and coal prices, upstream methane emissions from shale gas, capital costs for nuclear and coal with carbon capture and storage (CCS), public acceptance of carbon dioxide (CO_2) storage, and performance of gas-turbine-based technologies will influence investment dynamics through 2050. It explores how these long-run uncertainties can impact

[7]The focus on the electric power sector in the capacity planning and R&D portfolio chapters reflects the extensive research and policy consideration devoted to the industry, which also makes data more easily accessible to construct the models.

[8]Some transmission facilities in the US are nearly a century old.

near-term investments and how hedging strategies can reduce exposure to downside losses while preserving options for upside gains from volatility.

Capacity planning in the power sector fits the stochastic programming paradigm in that strategies are modified over time in light of new information about policies, technologies, and resources. The optimization problem is framed from the perspective of utilities and generators, where the objective function is to minimize discounted energy system costs. The primary decision objectives are to determine what types of generating capacity to build and operate and when such units should come online. The goal of long-term planning for utilities is to ensure that adequate resources are available to reliably serve demand while balancing other objectives of shareholders, ratepayers, and the general public (as well as other system constraints).

This research characterizes uncertainty through a range of approaches, including statistical analyses, expert elicitations, and econometric modeling. It is among the first to use a stochastic programming framework in a large-scale energy-economic model with a wide range of simultaneous uncertainties and many scenarios. The model is also the first to incorporate upstream emissions from shale gas production into an energy-economic model that can examine tradeoffs between uncertain life-cycle costs and environmental impacts of different technologies.

1.3.2 Energy Technology R&D Portfolio Management

Beyond issues of capacity planning, the power sector also contends with the uncertain availability and performance of technologies, which is influenced by public and private R&D investment decisions. Additionally, under conditions of deep uncertainty about climate change, the optimal strategy may not only be a single action like mitigation but also may involve building the capacity to respond to uncertainty in future periods through investments like R&D. Although it is a technology-based sector, electric utilities do not invest heavily in R&D relative to other US industries. Current R&D spending by investor-owned utilities is only about 0.3 percent of revenues, which is much lower than the 2.8 percent allocation of the total economy toward R&D expenditures (Anadon et al., 2011).

Given the importance of managing technological change for industry, government, and society, the research in Chapter 6 formulates and employs a novel R&D portfolio allocation framework. The objective of this research is to use stochastic programming metrics and tools to inform questions of energy technology R&D strategy.

Uncertainty is a fundamental characteristic of the R&D process. The stochastic and dynamic aspects of these questions are significant structural features of R&D strategy, including uncertainty in market and policy conditions, the relationship between R&D investment and technological outcomes, and the ability to adjust decisions over time based on learning. This research makes suggestions about how uncertainty should be represented in the R&D portfolio problem and adds many innovations to the literature.

This chapter informs questions about how to value technological advances. Measuring the benefits of R&D expenditures typically comes through two steps—namely, the relationship between R&D portfolio investments and potential outcomes as well as the valuation of these outcomes. This research investigates R&D success valuations in a sequential decision-making setting, and a novel contribution of the model is its stochastic diffusion mapping through the two-stage capacity planning model described in Chapter 3. The key attribute of the model is that, when R&D funding and first-stage decisions are made, the realizations of other exogenous uncertainties (e.g., abatement stringency) are unknown, which translates into uncertainty about diffusion markets for technologies upon which the R&D acts. This approach provides a more accurate representation of the R&D decision-maker's dilemma in which allocation decisions must be made in an uncertain market environment, where prospective conditions are subject to many contemporaneous sources of uncertainty.

Another objective of this work is to parameterize innovation production functions using results derived from expert elicitations rather than using *ad-hoc* values. Previous innovation production function analyses use stylized values that are the same across all technologies. In contrast, the work here provides some empirical grounding for the chosen values that link model representations with on-the-ground expectations for R&D program characteristics. This trait, combined with the stochastic valuation model and larger portfolio of R&D programs, suggests that model outputs offer a greater degree of normative decision support compared with previous analyses.

An overarching goal of this chapter is to embed the treatment of R&D control within the broader context of energy modeling. There is a long history of using energy-economic and integrated assessment models (IAMs) to quantify the benefits of technological developments but only a limited amount of work that explicitly links these models with an R&D portfolio framework. The research here leverages the experience, tools, and insights from the energy modeling community to better understand the relationship between R&D investment decisions and technological change, market diffusion, and environmental outcomes.

The primary research questions of the R&D portfolio management work include:

- Using the best-available expert elicitations for future technological states and up-to-date characterizations for other uncertainties, what R&D portfolio strategies will maximize expected welfare?
- How do recommended investment strategies vary based on the assumed decision-making approach?
- What are the impacts of R&D on private and societal costs as well as on environmental outcomes?

1.3.3 Fat-Tailed Uncertainty, Learning, and Climate Policy

The presence of uncertainty is common in energy and environmental decision-making, especially in formulating policy responses to climate change. Deep structural uncertainty about the climate system combined with tremendous challenges associated with quantifying the economic impacts of severe climate change pose many interrelated conceptual, methodological, and ethical difficulties. In the energy modeling community, "among the subjects that deserve further in-depth investigation, the issue of uncertainty emerges as, perhaps, the most prominent" (Haurie, Tavoni, and van der Zwaan, 2012), especially for the economic analysis of climate change. Problems of global change require more robust tools to characterize and analyze uncertainty than conventional approaches offer (Morgan et al., 1999). Although insuring against low-probability, high-impact climate risks has been a leading justification for mitigation for many decades (Manne and Richels, 1993), recent developments have refocused attention on how fat-tailed uncertainty, where probabilities of rare events decline relatively slowly in the upper tail of a distribution, may influence the urgency and degree of precautionary action that is warranted for abatement measures.

Like the other applications in this dissertation, the recognition of the importance of uncertainty in climate change economics has outpaced its implementation in models. IAMs are largely deterministic and focus on expected-value forecasts, using one-way sensitivities to assess the impacts of uncertain or contentious parameters.[9] In the rare instances when uncertainty is more formally incorporated into the analysis, uncertainty propagation techniques are typically used with thin-tailed probability density functions (Nordhaus, 2008; Hope, 2006). These approaches implicitly assume perfect information for each simulation run (i.e., implying a learn-then-act approach in which the uncertain state is revealed before decisions are made). This characteristic means that such *ex-post* approaches, while analytically simple, cannot offer guidance in determining *ex-ante* hedging strategies. Sequential decision-making frameworks incorporate uncertainty explicitly and address limitations of foresight by determining optimal policies in multiple stages based on updated information.

The objective of Chapter 7 is to examine how sequential decision-making frameworks, concepts, and metrics can be used to inform risk management in climate policy. Specifically, this work examines the impact of fat-tailed uncertainty about the climate sensitivity parameter and the potential for learning on optimal near-term abatement. Computational experiments using this stochastic decision model answer questions about whether policy recommendations are robust to the specifications of distributions, damages, and discounting. The focus is on the timing and stringency of global climate policy and not on questions of mechanism design for coordinating international efforts to control emissions.

[9] Another disadvantage of using expected-value or best-estimate forecasts for decision-making is that such approaches may "inhibit deliberations among individuals with differing expectations and values, because by definition best-estimate predictions privilege some expectations and values over others" (Lempert, Groves, and Fischbach, 2013).

Another goal is to quantify the value of learning and midcourse corrections on reducing consumption risks imposed by uncertain damages from climate change. In the presence of strong stock-accumulation inertias and sunk emissions, Weitzman (2012) assumes that lags involved in climate system preclude the ability to learn about catastrophic impacts until damages arrive. Other authors (Nordhaus, 2012; Kousky et al., 2010; Nordhaus, 2009; Yohe and Tol, 2007) have criticized Weitzman's assumptions as being unrealistically pessimistic and suggest that the pace of climate change will allow for the possibility of learning over time, deploying negative emissions technologies, or using emergency geoengineering if warming is unexpectedly high. Despite these criticisms, no work has investigated the degree to which assumptions about learning influence the conclusions of the model in Weitzman (2012). This research examines the influence of different learning rates on near-term abatement decisions and discusses the normative implications of considering uncertainty about the evolution of uncertainty over time.

A high-level motivation for this work is to investigate to what degree policy prescriptions in IAMs are robust to conventional assumptions about thin-tailed probabilities, perfect foresight, and quadratic damages. Forecasts for the benefits and costs of abatement are predicated on many assumptions, including the response of the climate system to emissions-driven forcing, severity of economic impacts of climate change, rate of invention and innovation for abatement technologies, growth rates of emerging economies, discounting assumptions, and many other factors that are difficult to conceptualize let alone to predict. Since there are many plausible ways to select model parameters related to scientific knowledge, human behavior, and value judgments, decision-makers have a difficult time adjudicating between alternatives when there is little basis for agreement between studies, as stakeholders can rationalize a set of assumptions to cohere with policies they support on the basis of cultural, financial, or other considerations.

The research in Chapter 7 provides a unified framework for isolating these effects and for showing how policy guidance from IAMs may rely strongly on these assumptions. The sequential decision-making approach used here reconfigures the policy debate by considering a wider set of potential outcomes and adaptive actions for risk management. The results do not propose specific and definitive solutions to these complex issues but draw general insights, which are useful starting points in understanding potential modeling limitations and refocusing research attention on problems related to uncertainty and learning. Ultimately, such integrated assessments under uncertainty can offer guidance in thinking through bounding questions about the importance of climate change.

1.4 Dissertation Organization

The three applications discussed in Section 1.3 are linked by the mathematical framework of stochastic programming. Chapter 2 begins the body of the dissertation by discussing the tradeoffs, benefits, and shortcomings of various approaches to modeling uncertainty and motivates the appropriateness of stochastic programming for the applications in this dissertation. In addition to defining many of the terms used throughout the dissertation and formulating metrics for evaluating uncertainty, the chapter reviews literature related to uncertainty analysis approaches and their applications in energy modeling with a focus on electric sector capacity planning.

Chapters 3 through 5 comprise the electric sector capacity planning work. Chapter 3 formulates a two-stage stochastic programming model of capacity planning and dispatch for the US electric power sector. Chapter 4 presents model results for the dynamics of capacity planning under a range of technological, economic, and policy-related uncertainties. Chapter 5 analyzes the capacity planning results in the broader context of uncertainty, learning, irreversibility, and optionality in the power sector.

Chapter 6 introduces a novel stochastic R&D portfolio management framework and presents results for energy technology R&D strategy in a carbon-constrained world. Chapter 7 examines the impact of fat-tailed uncertainty about climate change on near-term abatement decisions using a sequential decision-making framework. Finally, Chapter 8 offers overarching conclusions and suggests fruitful directions for future research.

The appendices present additional information on important areas of the dissertation where inclusion in the body would impede the flow of the narrative. Appendix A discusses details of an approach for estimating correlated probability distributions for natural gas and coal prices, which is used throughout the capacity planning work. Appendix B examines the factors that enhance the reliability of expert elicitations and discusses unresolved questions about best practices for elicitation protocols. These insights are applied in a case study to understand the current state of knowledge regarding the future of gas turbine systems for electricity generation, which is used as an input to the capacity planning model. Appendix C uses a stylized example to demonstrate how the value of R&D success can vary based on the decision-making approach used for capacity planning decisions.

Chapter 2

Methodological Introduction and Literature Review

The purpose of this chapter is to situate the modeling framework of stochastic programming within the broader terrain of uncertainty analysis. This chapter describes the tradeoffs, benefits, and shortcomings of various approaches to modeling uncertainty and motivates the appropriateness of stochastic programming for the applications in this thesis.

The chapter begins with an introduction to the stochastic programming paradigm and describes its appeal for energy modeling. Sections 2.2 and 2.3 compare stochastic programming with other forms of uncertainty analysis and highlight its strengths over other sequential decision-making approaches. Next, Section 2.4 provides definitions associated with the stochastic programming framework, which will be used in subsequent chapters of this dissertation. After these discussions of uncertainty analysis approaches and their applications in energy modeling, Section 2.5 reviews related literature in electric sector capacity planning. This section discusses previous publications and assesses the novel contributions of the research in Chapters 3–5.

2.1 Stochastic Programming Background

Stochastic programming is a framework for solving decision problems under uncertainty. Determining optimal strategies requires the consideration of unknown parameter values, which are features of nearly all decision problems. This approach generates more robust decisions than deterministic planning.[1] The resulting strategies perform reasonably well across a wide range of plausible realizations, reflecting both a normative choice framework and a descriptive criterion that many decision-makers use under uncertainty (March, 1994).

Like all sequential decision-making approaches (Kann and Weyant, 2000), stochastic programming models find optimal policies during multiple stages over time. These actions are selected given joint probability distributions over a set of potential outcomes in subsequent stages.[2] This structure captures the dynamic nature of decision problems where policies are revised as new information becomes available, as in the cases of capacity planning, induced technical change, and climate policy.

The objective of stochastic programming is to develop hedging strategies. Although these solutions are unlikely to be ideally suited for every state of the world, these strategies are optimal *ex ante* and provide adaptability and resilience under a variety of future scenarios.[3] When uncertainty is prevalent and the long-term consequences of near-term decisions are imperfectly understood, resilience allows decision-makers to shape available options, to cope with the unknown, to learn from errors, and to exploit new information as conditions change. Thus, the stochastic programming framework strikes a balance between optimality and robustness.[4]

The stochastic programming framework is best suited for decision problems that

[1] Here, robustness is operationalized by minimizing the probability-weighted sum of discounted costs across all possible scenarios.

[2] This work adopts a subjectivist Bayesian interpretation of probability in which the probability of an event is a measure of the degree of belief that the event will occur given all relevant and available information (Morgan and Henrion, 1990).

[3] This behavior gives rise to the notion that the early availability of information can improve performance, which is the basis for the expected value of perfect information in Chapter 3.3.3.

[4] Optimality-based approaches can be vulnerable to overconfidence but are simpler to construct and to understand, whereas robust approaches require more analysis yet give a more complete portrait of risk.

satisfy a few specific assumptions, as described in greater detail in Section 2.2. The formulation of a two-stage stochastic program assumes that:

- Known probability distributions can be attached to uncertain parameter values
- The resolution time (i.e., the period when updated information will be available) is known at the beginning of the time horizon
- The timing of parametric uncertainty resolution is independent of decisions
- All functions are measurable with respect to the sigma-algebra associated with the probability space of the problem and that all expectations exist (Dupačová, Hurt, and Štěpán, 2002)

The standard approach to decision problems under uncertainty is to replace random parameters by their expected values (or another measure of central tendency) and to solve the resulting deterministic optimization problem. However, this expected-value strategy is optimized only for a specific set of conditions, and when considering the range of possible futures that may come to pass, this solution may be far from optimal. Disregarding inherently random characteristics of decision problems may limit the usefulness of the resulting solutions, which are likely brittle. Sam Savage refers to this notion that "plans based on *average* assumptions are wrong on *average*" as the "flaw of averages" (Savage, 2009).

For energy-related decisions, if planners assume mean values and encounter something unexpected, they risk imposing economic and environmental burdens on a range of stakeholders. For instance, Ho Chi Minh City responded to their vulnerability to routine flooding by undertaking a series of multi-billion dollar infrastructure projects over the past 15 years. Encouraged to take action by rising sea levels and increased precipitation caused by climate change, the city viewed these investments as important adaptation and risk management tools to cope with an uncertain climate and with development in low-lying areas. City developers designed canals, pipes, dikes, river barriers, and gates based on best-guess estimates for future climate change and carbon dioxide (CO_2) emissions. However, climate impacts have exceeded forecasts in the years since planning began, as some variables are already higher than the distant-future design specifications. These surprises were especially damaging due to increased urbanization in low-density areas, which was partially bolstered by the illusion of safety provided by the infrastructure investments. Consequently, these flood risk mitigation decisions proved brittle in the face of misspecified predictions and actually increased the vulnerability of Ho Chi Minh City (Hallegatte et al., 2012).

A second approach to uncertainty analysis is to consider all possible future scenarios and to design decisions that are optimal under each state of the world. Although this approach performs well when uncertainty can be resolved early or when opportunity costs of delay are low, such decision contexts are rare. Hence, this wait-and-see approach is not implementable.

A related alternative is to use a heuristic procedure to devise a composite strategy from multiple wait-and-see solutions. This scenario analysis approach is often coupled with sensitivity analyses to determine the robustness of the resulting solution. However, this technique is susceptible to the same shortcomings as the expected-value approach—namely, that the solution may be far from optimal in an expected-utility sense. These scenario-based approaches are especially inadequate from a decision-making perspective. Having a multiplicity of strategies that are tailored to specific states of the world leaves decision-makers in a quandary about how to choose among alternatives and to develop near-term strategies. The existence of many contradictory scenarios fails to provide unambiguous policy-relevant insights, especially when no probabilities are attached to these scenarios. Furthermore, the unknown cost of uncertainty makes it challenging to rank different uncertainties on the basis of their importance, which is a critical facet of policy analysis and model construction.

Finally, the stochastic programming (or, more generally, sequential decision-making) framework remedies these shortcomings by identifying a single hedging strategy, which provides clearer guidance to decision-makers. This approach balances, *inter alia*, the risk of waiting to learn more information with the risk of premature action. It accounts for a wider range of probability-weighted outcomes by striking a reasonable compromise between expected performance and risks from manifold sources of prediction error. Hedging is appropriate if near-term decisions (i.e., those made before uncertainties are resolved) are different from expected-value solutions.[5]

Stochastic strategies create contingency plans that respond to opportunities and pitfalls through their adaptive framework of midcourse corrections. In this sense, stochastic programming develops "antifragile strategies" (Taleb, 2012) that not only reduce the possibility of downside losses but also provide options for upside gains from volatility.[6] This distinction is especially important in the context of energy systems where decision-makers must contend with many simultaneous uncertainties, including both negative Black Swans (e.g., devastating natural disasters like Hurricane Sandy in October 2012 and the Tōhoku earthquake and tsunami in March 2011) and positive ones (e.g., abundant and economical shale gas reserves). Due to the ineluctability and pervasiveness of prediction error, the objective of adaptive approaches is to develop strategies that limit and localize impacts of mistakes and that do not rely on accurate, agreed-upon forecasts.

There are many factors contributing to the recent boom in applications of stochastic programming methods, especially in non-energy fields (Birge and Louveaux, 2011; Wallace and Ziemba, 2005). The primary factors are advances in computing technology as well as in computational and analytical methods for optimizing mathematical programming models (e.g., decomposition techniques). More readily available commercial solvers have increased the usability and accessibility of these methods. For instance, this research uses the DECIS system for solving large-scale stochastic programming problems, which can employ sophisticated solution strategies while interfacing with the widely used GAMS algebraic modeling system (Infanger, 1999). It is fortuitous that these computing resources and analytical methods are emerging concurrently with long-range policy issues related to economic development, the environment, and public safety, which require such methods to solve decision problems

[5]As described in Section 2.4, stochastic programming methods are most useful when the resulting strategies are not identical to any wait-and-see deterministic strategies. This characteristic indicates that optimal energy system and technology research and development decisions are not easily discernible without tools that treat uncertainty explicitly.

[6]Taleb contrasts fragile strategies that leave decision-makers "*at best* unharmed" and robust strategies that leave them "*at best* and *at worst* unharmed" with antifragile strategies that leave decision-makers "*at worst* unharmed" (Taleb, 2012).

under pervasive uncertainty.[7]

2.2 Comparison of Uncertainty Analysis Approaches

2.2.1 Sensitivity and Scenario Analysis

The two most basic and common treatments of uncertainty are sensitivity analysis and scenario analysis. *Sensitivity analysis* assesses the sensitivity of model outputs to changes in inputs. This approach is commonly implemented by varying uncertain parameters across their permissible ranges while holding all other values constant and then determining the direction and magnitude of change for outputs of interest.[8] These perturbations can give modelers and decision-makers a sense of potential variation in outputs when input parameters are misestimated.

Sensitivity analysis can be useful in identifying parameters that have the largest effect on results and in finding break-even points where optimal decisions change. For modelers, sensitivity analysis can suggest parameters to treat stochastically and is a prerequisite to more sophisticated forms of uncertainty analysis.[9] This approach is advantageous due to its intuitive appeal, synergies with model diagnostics, and straightforward implementation, as it can be performed with little extra computational complexity or model modifications.

Despite these attractive features, sensitivity analysis suffers from many limitations. The range of values may not reveal the uncertainty involved, especially when

[7]To illustrate progress in these areas over the past two decades, consider the following passage from William Nordhaus' *Managing the Global Commons* (Nordhaus, 1994), "Programming the DICE model for five [states of the world] on a PC is moderately straightforward, and solutions on a top-of-the-line PC take thirty minutes as of summer 1993. Clearly, solving for even a few dozen of the 400,000 [states of the world] would require supercomputers." This instantiation of DICE had roughly 4,000 nonzero coefficients in the deterministic problem matrix. In contrast, the model described in Chapter 3 has over 180,000 nonzero elements, 10,000 scenarios for some runs, and takes under 30 minutes to solve on a far-from-state-of-the-art personal computer.

[8]This approach also can be operationalized by taking the partial derivative of the model output with respect to an input (Katz, 2002).

[9]As described in Chapter 3.3.3, sensitivity analysis is complementary to stochastic programming in establishing first-order bounds on metrics like the value of the stochastic solution, which can guide the selection of uncertainties.

random variables are correlated and when models are nonlinear. One-way sensitivities are typically used for model diagnostics even though interactions between parameters are frequently important. Since most methods are based on linear approximations, their conclusions are typically only valid locally and are not suited for highly nonlinear systems. Finally, sensitivity analyses are often predicated on the assumption that the model structure is correct.

A related treatment of uncertainty is through *scenario analysis*. Scenario analysis compares outputs of model runs with different combinations of assumptions and parameters. The climate community frequently employs scenario analysis (IPCC, 2000) through the intuitive logics school, which emphasizes informative scenarios that identify key decision-relevant uncertainties (Lempert, 2012). Like sensitivity analysis, scenario analysis has a few intuitive and pedagogical advantages, which stem from its straightforward implementation and low computational effort.

The shortcomings of scenario analysis make it poorly suited for performing rigorous uncertainty analysis. Scenarios are generally not weighted by probabilities, which makes it challenging to navigate through a multiplicity of feasible scenarios. Ambiguity from the absence of likelihoods leaves decision-makers more susceptible to being swayed toward a favored alternative, which may be very different from the strategy suggested by a probability-weighted decision framework (Schneider and Kuntz-Duriseti, 2002).[10] Additionally, there is concern that scenario analysis may lead to overconfidence and underestimation of uncertainty (Morgan, 2011). As more detail is added to a scenario to enhance its clarity, the availability heuristic suggests that the verisimilitude of that particular state of the world will increase, which consequently will increase the outcome's perceived likelihood while making it more difficult to imagine other plausible and equally likely outcomes. Furthermore, given that there are no best practices for scenario design, scenario analysis methods can omit relevant surprises and discontinuities associated with low-probability, high-impact outcomes. These omissions may be especially problematic when scenario design focuses on creating a few detailed, self-consistent storylines (Postma and Liebl, 2005).

[10] The scenario literature offers supporting evidence for this claim. Many users of the IPCC's Special Report Emissions Scenarios use *ad-hoc* rules for assigning probabilities (since no likelihoods are given) such as assuming equal probabilities for all scenarios (Parson et al., 2007).

2.2.2 Uncertainty Propagation and Sequential Decision-Making

There are many different approaches to incorporating uncertainty analysis into large-scale models. Most techniques can be classified as either uncertainty propagation or sequential decision-making under uncertainty (Kann and Weyant, 2000).

The more common approach is *uncertainty propagation*, which involves sampling from distributions for uncertain parameters, propagating them through a deterministic model, and creating output distributions. The simplest implementations of this approach are Monte Carlo methods, which rely on repeated sampling from joint distributions over uncertain parameters and then propagating uncertainty.[11] This method generates probabilities over outputs of interest or risk profiles for given strategies. Such outputs provide decision-makers with a sense of risk, which can be important in contexts with nonlinearities and risk aversion. Uncertainty propagation does not incorporate learning, since the optimal strategy is only determined once at the beginning of the time horizon.

Although uncertainty propagation techniques can be used with minimal model modification, making them easy to use with existing deterministic models, these approaches have a few shortcomings:

- They do not give adequate guidance to decision-makers about selecting strategies before the uncertain state of the world is revealed. Propagation assumes forward-looking perfect information across the time horizon for each simulation run (i.e., implying a learn-then-act approach in which the uncertain state is revealed before decisions are made). This characteristic makes it difficult to determine near-term hedging strategies, especially when many uncertainties are considered simultaneously.

- Parameters may contribute to uncertainty but may be irrelevant to decisions, which is captured by value of information calculations.[12] Thus, uncertainty propagation many not be adequate for identifying policy-relevant parameters.

[11] Efficient sampling schemes are frequently used with these simulations to lessen the computational burden associated with complex joint distributions.

[12] In other words, the objective function value can vary greatly with changes to input parameters, but the optimal decision variables may be the same.

- Like sequential decision-making, propagation approaches require the specification of joint distributions for all uncertain random variables. Even if model builders are unbiased in characterizing these distributions, uncertainty quantification can be challenging when correlations are present.

Ultimately, propagation approaches make it difficult to translate risk profiles into actionable insights. These techniques do not incorporate the flexibility of revising strategies as more information becomes available and realized parameters differ from their initial expectations. The dimension of adaptability to surprise is not captured through uncertainty propagation.

In contrast, *sequential decision-making* frameworks like stochastic programming identify hedging strategies that balance the risks of premature action with those of delay given the decision-maker's current state of information.[13] These adaptive decision models address limitations of foresight by determining optimal policies in multiple stages based on updated information. This characteristic offers a more robust treatment of uncertainty than propagation models, which select policies once and do not incorporate learning thereafter.[14] The two-stage stochastic programming framework describes the problem many decision-makers face of identifying optimal short-term strategies in face of long-term uncertainty while accounting for how learning and adaptability may impact optimal policies.

Although sequential decision-making frameworks can provide decision support under uncertainty, such approaches have a few drawbacks:

- These techniques have high computational burdens and costs of complexity compared with other approaches to uncertainty analysis.[15] The curse of dimensionality refers to the multiplicative model growth with the number of stages and number of scenarios in each stage.

[13]Uncertainty propagation approaches can be complementary if simulations are used to generate scenarios, which can then be used as inputs to a stochastic programming framework.

[14]Note that, although a decision-maker's information may increase over time, uncertainty does not necessarily decrease (Hannart, Ghil, and Dufresne, 2013), as discussed in Chapter 7.

[15]There is also typically a tradeoff between the fidelity to a decision-maker's problem and the ease of interpreting the results to extract policy-relevant insights.

- It is more difficult to convert an existing deterministic model into a sequential one compared with uncertainty propagation approaches. These changes require structural modifications. The computational complexity of these more sophisticated optimization models entails a tradeoff between the resolution of the modeled system and the exhaustiveness of the characterization of uncertainty.
- The learning process is not typically straightforward due to factors like noise, measurement error, stochasticity, and imperfect knowledge of system dynamics.

Like uncertainty propagation, sequential decision-making approaches focus on parametric uncertainty and dodge cumbersome issues of structural uncertainty. Additionally, the complexity of the problem may mean that model construction is delegated to specialists, which creates a danger that the decision-maker's understanding of and commitment to the model results may be lower due to reduced transparency.

2.2.3 Decision-Making Approaches under Deep Uncertainty

Conventional probabilistic approaches characterize uncertainty through joint probability density functions; however, these frameworks do not address deep uncertainty, imprecision, or ambiguity in these estimates.[16] For decision-making problems under opacity, there are many ways to describe conditions of deep uncertainty and alternative decision criteria to managing risk. These frameworks are useful in information-sparse situations with extended time horizons and long-term impacts (Lempert, Popper, and Bankes, 2003).

When confronted with deep uncertainty, there are a few methods to treat uncertainty that do not rely on a single joint probability density function. Axiomatic approaches like belief functions and fuzzy logic offer alternative frameworks that relax probability axioms to represent the notion that "one can gain or lose confidence in one of a mutually exclusive set of events without necessarily gaining or losing confidence in the other events" (Morgan et al., 2009). A second method for representing imprecision is through alternative sets, which may either use a range of probabilities

[16] Knight (1921) famously differentiates between risk (with well-quantified and knowable probabilities) and uncertainty (with poorly understood and difficult-to-know probabilities).

or non-probabilistic sets of potential scenarios. Finally, scenario-based descriptions of uncertainty present alternative states of the world but focus on characterizations that engage the imaginative capacity of decision-makers (Lempert and McKay, 2011).

The literature formulates and applies many different decision criteria under conditions of deep uncertainty. Robust optimization can be either probabilistic or non-probabilistic but shares the common characteristic of exchanging some degree of optimal performance for reduced exposure to risk, much like stochastic programming. These frameworks focus on robustness over a range of potential conditions and aim to "illuminate the vulnerabilities of proposed policies" (Lempert, 2012). These characteristics make robust optimization especially advantageous when stakeholders cannot agree on functional relationships, probability distributions, or valuing outcomes.

Two of the most common non-probabilistic decision criteria are minimax regret and maximin.[17] The minimax-regret approach minimizes the maximum or worst-case regret, where regret is the difference between the payoff of the best policy for a specific scenario and the actual payoff. The maximin criterion ranks policies by their worst-case outcomes and does not incorporate regret. Maximin focuses solely on downside risk, which results in more conservative policy recommendations compared with minimax regret. Despite their analytical simplicity, these approaches provide incomplete measures of the consequences of suboptimal planning decisions. For applications in energy and climate change, there are many plausible values in the supports of probability distributions associated with low-probability, high-impact scenarios that would lead to policy recommendations of exchanging an excessive degree of expected returns for lower risk.[18] Recent research bridges expected utility theory with minimax approaches using value-at-risk criteria from the finance literature, which determines optimal strategies under probabilistically defined uncertainty given satisficing constraints regarding outcomes in the tails of distributions.

[17] The premise that some problems are too information-sparse for probabilistic analysis is controversial (Bier et al., 1999).

[18] There are many well-known pathologies associated with the minimax-regret and maximin approaches, especially regarding highly unlikely worst-case scenarios (Savage, 1954).

2.3 Sequential Decision-Making Approaches and Energy Modeling

2.3.1 Adaptive Strategies

The guiding principle behind sequential decision-making is that adaptive systems should account for future learning, potential for surprises, capabilities of stakeholders, and consequences of suboptimal decisions. As suggested in Section 2.2, sequential decision-making approaches share the common feature of revising strategies over time as new information becomes available. This structure requires multiple decision points that are temporally separated, allowing a decision-maker's information about uncertainties to change over time and making midcourse corrections possible.

The analytic need for sequential decision-making approaches in energy, economic, and environmental modeling is due to a combination of uncertainty and inertia in these systems. In the absence of uncertainty, decision-makers could formulate an optimal response trajectory for the entire planning horizon, which obviates the need for iterative policies or midcourse corrections. Without inertia, optimal near-term decisions in any period would depend only on the conditions, information, and uncertainty observed in that particular period, which suggests that corrective actions would be unnecessary. However, when both uncertainty and inertia are present, strategies may be revealed to be suboptimal, and adjustment decisions must be made in response to these changing conditions.

Sequential decision-making models provide a more realistic depiction of problems related to capacity planning, technology research and development (R&D) management, and climate policy than their deterministic, perfect-foresight counterparts. The sequential paradigm involves anticipation, learning, and adaptability, where decisions are modified over time in light of new observations. In the context of climate policy, this ability to redirect decisions is an important component of effective and efficient policies, as the experiences with early implementation provide valuable information on compliance costs, response of the climate system, and adaptive capacity. Rather than a “large, one-shot, ‘bet-the-planet’ decision” (Weyant, 2008), sequential frameworks incorporate flexibility into planning decisions, which helps agents to enhance adaptability and to limit vulnerability.

With very few exceptions, large-scale energy-economic models that are used to inform energy and climate decisions do not examine the effects of sequential strategies. Although there is wide recognition that uncertainty plays a critical role in these domains, the explicit treatment of uncertainty and sequential decision-making has been restricted.[19] The intertemporal optimization structures found in many integrated assessment models do not incorporate multiple decision stages, which consequently means that they do not include sequential decision-making structures (Parson and Karwat, 2011). Myopic frameworks make static optimization decisions during each time period in response to evolving conditions but do not consider the impacts of current decisions on anticipated future benefits and costs apart from current rates of return. Most intertemporal optimization models identify a single decision strategy over the entire time horizon assuming perfect foresight, which specifies the complete decision strategy by the forward-looking representative agent and does not account for the possibility of midcourse adjustments.

Uncertainty analysis in energy modeling typically involves sensitivity analysis or uncertainty propagation, where exogenous probability distributions are propagated through deterministic models. These techniques do not contain the learning characteristics of sequential decision-making. The curse of dimensionality makes it challenging to solve stochastic optimization problems, but these issues are intensified in the domain of large-scale integrated assessment models, which are often nonlinear as well. From a formulation perspective, sequential models require the specification of information structures describing the evolution of uncertainties over time and of intertemporal structures of decisions (with the potential for dynamic choice sets).

[19] The Intergovernmental Panel on Climate Change’s Fourth Assessment Report states, “Responding to climate change involves an iterative risk management process...” (IPCC, 2007).

2.3.2 Sequential Decision-Making Approaches

Given the value and need for sequential decision-making frameworks in energy-economic modeling, it is important to adjudicate on the appropriateness of different mathematical approaches in this context.

The two most common frameworks for modeling and solving sequential decision problems are stochastic programming and dynamic programming.[20] The concepts, formulations, and solution algorithms for these two approaches were independently developed at the same time (Dupačová and Sladký, 2002), but despite these similarities, there has been little acknowledgement of the parallels and complementarities between these sequential decision-making approaches.

Multi-stage stochastic programs are comparable to discrete-time stochastic dynamic programming problems in that they deal with dynamic and stochastic decisions (Dupačová, Hurt, and Štěpán, 2002). The primary distinction between these approaches is the solution concept, implying that selection of an appropriate framework is driven by the structure of problem, available data, and tools for solving the resulting model. Stochastic programs typically involve problems with many decision variables (with many potential values) but a limited number of stages. Dynamic programming is most useful when problems have large numbers of decision stages but limited state spaces or when uncertainties resolve at different or unknown times. Thus, the main emphasis for stochastic programs is typically on insights about first-stage hedging decisions, whereas the primary interest in dynamic programming is the decision rule itself.

Both approaches also suffer from the curse of dimensionality. For dynamic programs, the curse results from the number of states and dimensionality of the state and control spaces. For stochastic programs, the curse of dimensionality is due to the number of stages and scenarios.

[20] Other related mathematical frameworks include decision analysis, stochastic control theory, optimization of discrete-event simulations, and Markov decision processes (Shapiro, Dentcheva, and Ruszczyński, 2009).

2.3.3 Benefits and Drawbacks of Stochastic Programming

In the domains of climate change policy and technology R&D strategy, the two-stage stochastic programming approach strikes a balance between model fidelity and tractability for formulating and solving problems. A widespread misconception is that stochastic programming is only useful in modeling contexts with very few uncertainties, and this misunderstanding has likely limited the applications of this approach. Although this claim may be true of solving deterministic equivalent problems directly, one advantage of stochastic programming is its ability to use different solution strategies and to exploit powerful solvers that employ techniques like Benders decomposition and Monte Carlo sampling with variance reduction techniques (Glynn and Infanger, 2013; Infanger, 1994). Many integrated assessment and energy-economic models are known for the high dimensionality needed to represent technology competition between a wide array of supply- and demand-side technologies (that often have considerable cost and performance detail), energy carriers, and regional characteristics. This high dimensionality makes technology-rich models well-suited for stochastic programming settings.

The two-stage stochastic programming approach also has distinct benefits for model formulation, particularly for energy models. Many uncertainties for energy systems and climate change (e.g., the timing and stringency of climate policy, technological characteristics) can be described as long-term, low-frequency uncertainties (Ryan, McCalley, and Woodruff, 2011). These uncertainties may not occur repeatedly, and analysts cannot reliably use historical time series data, functional relationships, or relative frequencies to construct probability distributions for the future.[21] For example, technological breakthroughs are inherently unique, particularly for nascent technologies, which means that historical data for carbon capture and storage (CCS) analogs like flue-gas desulfurization may give little relevant guidance for distributions over future CCS costs (Rai, Victor, and Thurber, 2010). Thus, when past data are

[21] As pointed out by Weitzman (2009), uncertainties like the climate sensitivity parameter inherently have diffuse distributions. Future realizations of parameter values, particularly those outside of the range of experience, are not adequately covered in past observations, which makes it challenging to learn limiting tail behavior through induction using finite historical samples. Chapter 7 discusses modeling and policy questions about fat-tailed uncertainty in the context of climate policy.

unavailable or of limited use, modelers often turn to the structured process of expert elicitations to encode probabilities of individuals with expertise.[22]

As discussed in Appendix B, expert elicitations for energy technology cost, performance, and R&D have been undertaken in recent years, and research efforts to integrate and communicate these results suggest that the frequency and detail of these elicitations will increase in future years. There is extensive experimental literature suggesting that both laypeople and experts are poor at dealing with correlational structures (Morgan and Henrion, 1990), which means that uncertainties in elicitations are often treated as independent random variables.[23] In light of this limitation, an appropriate modeling framework for this class of problem should have as few conditional probabilities as possible (but no fewer). Thus, the formulation of two-stage stochastic programs makes this framework attractive for energy-economic modeling, especially since dynamic programming requires the specification of many conditional transition probabilities.[24] The fact that stochastic programming interfaces naturally with expert elicitations is a strong benefit of this framework and makes it well-suited for future research.

The stochastic programming framework has many other benefits in areas of model formulation, implementation, diagnostics, and interpretation. In calculating metrics for evaluating the importance of various uncertainties, this approach forces modelers to "reflect more on the appropriateness of... model formulations and parameter assumptions than might otherwise be the case" (Weyant, 2008) and to make use of long-neglected diagnostic tools like tornado diagrams. For instance, in order to calculate the expected value of perfect information (EVPI), a modeler must run perfect

[22]Appendix B discusses energy technology expert elicitations in greater detail and presents an application for natural gas turbine efficiencies.

[23]In order to retain some degree of probabilistic dependence while avoiding pitfalls of conditional elicitations, it is preferable to circumvent the issue by explicitly modeling the cause of the dependency. For instance, if future costs of nuclear power plants are correlated with CCS-equipped coal facilities, this dependency may be caused by construction cost inflation, which can be incorporated in the model as an extra parameter.

[24]Although the dynamic programming setting is more appropriate when decision-dependent uncertainties play prominent roles, stochastic programs can be formulated to accommodate such structures (Baker and Solak, 2013).

foresight runs for each possible scenario (i.e., every combination of uncertain parameter values). This exercise implicitly constrains modelers to conduct a high-dimensional sensitivity analyses to stress test models instead of merely one-at-a-time sensitivities. Even outside of value of information calculations, these values can quantify how interactive parametric dependencies can influence model outputs (Butler et al., 2013). Additionally, these outputs provide valuable tools for organizing and prioritizing areas of inquiry, identifying interdependencies, and promoting a clearer understanding of the dynamics of these complex systems.

In summary, the stochastic programming framework is typically most fitting for problems with many decision variables, few stages, and discrete decision points, which are all characteristics of many energy-economic and integrated assessment models.

Although the stochastic programming approach offer many valuable features to decision-makers and modelers, this framework has a few drawbacks. First, stochastic programming and other sequential decision-making techniques have substantial computational burdens due to the curse of dimensionality (i.e., multiplicative model growth with the number of stages and scenarios). Although this methodological difficulty is extensively documented and can be mitigated in part through state-of-the-art algorithms (Infanger, 1994), the curse of dimensionality also applies to manipulating, interpreting, and presenting the results of sequential decision models. The burgeoning size of model outputs makes it challenging to extract decision-relevant insights as well as to understand and communicate model results. Second, marginal distributions are difficult and time-consuming to assess for many uncertainties, let alone joint distributions. Some model parameters may be recognized as uncertain though probabilities may be challenging to quantify, but other uncertainties may be unknowable or erroneously assumed to be certain. However, the difficulties associated with quantifying uncertainty may suggest that a two-stage stochastic programming approach represents a satisfactory balance between representing risk explicitly avoiding onerous correlational structures. Another counterpoint to this objection is that even crude approximations for distributions can aid decision-makers in creating strategic scenarios and formulating hedging strategies.[25]

[25]Parson and Karwat (2011) suggest that, in conditions of deep uncertainty replete with unknown

2.4 Stochastic Programming Terminology

The most studied and applied stochastic programming models are two-stage linear programs with recourse, which are used in this dissertation. In the two-stage problem, decision-makers can make reasonable estimates about the near-term future, but assumptions about the longer-term future are uncertain. These beliefs are quantified through probability distributions over second-stage outcomes. Although a more realistic approach would account for the potential for making decisions continuously as information is received, a two-stage approach is more tractable for formulating and solving. The framework also captures the delay induced by the decision-maker's difficulty in separating the true value associated with an uncertain system from limited and noisy system observations (i.e., distinguishing between short-term variability and long-term trends).

Uncertain parameters are represented as a set of scenarios or states of the world, which specifies both the full set of random variable realizations and their corresponding probabilities. This set contains a wide array of potential futures, which are assumed to be mutually exclusive and collectively exhaustive. Probabilities associated with specific scenarios can be assessed through simulations, theory, statistical analyses based on historical data, meta-analyses, and expert elicitations.

Stages are distinct from periods in stochastic programming. Periods are intervals in the time horizon. Stages are sets of consecutive periods that divide the time horizon based on realizations of uncertainties and information sets of decision-makers.

For two-stage stochastic programs, first-stage decisions (i.e., here-and-now commitments) occur before uncertainties are resolved. Solutions are deemed admissible if they satisfy all constraints for all scenarios.[26] This expected-utility framework assumes that the utility of an uncertain prospect is determined by taking the sum of

unknowns, decision-makers should acknowledge problems of overconfidence in the planning process and welcome the paradox of "expecting to be surprised."

[26] Non-admissible solutions may be valuable in contexts with constraints that may be violated in a feasible range without significant consequences (Ryan, McCalley, and Woodruff, 2011). Chance or probabilistic constraints in stochastic programs allow violations with specified probabilities (Birge and Louveaux, 2011).

the probability-weighted utilities of potential outcomes.[27]

The nonanticipativity assumption constrains decisions to be the same under all scenarios until the uncertainty resolution date.[28] This assumption formalizes the notions that decision-makers cannot accurately predict the state of the world *ex ante* and that decisions in the current stage do not depend on future realizations of random variables or on future decisions. Hence, decisions are nonanticipative of future outcomes and use only past information and beliefs about probabilities of future states. Implementable solutions satisfy the property that decision variables under all scenarios are indistinguishable during the first stage. In most applications, the first-stage hedging solution is the most valuable model output to decision-makers, as it informs near-term decisions and provides adaptive potential for making recourse decisions once the state of the world has been revealed.

Second-stage (recourse) decisions are made *ex post* after more information becomes available to the decision-maker. The second-stage deterministic optimization problem finds the optimal vector of decision variables after uncertain parameters are observed. The first-stage stochastic strategy is common to all outcomes, whereas a second-stage solution is specified for each realized scenario. These recourse decisions can be considered contingent actions that are taken only when the corresponding state of the world is realized. The stochastic programing framework assumes that scenario probabilities are independent of first-stage decisions.

Thus, the optimal stochastic strategy specifies a single first-stage decision vector and a collection of second-stage decisions, which defines the optimal recourse strategy under each possible outcome. Calculating the optimal hedging strategy requires the decision-maker to consider all potential outcomes of all uncertainties, the probabilities associated with individual outcomes, and the gains and losses of strategies under all scenarios (including recourse alternatives and irreversibilities). Flexibility is a key attribute of stochastic strategies and indicates the degree to which a system can

[27]Kahneman and Tversky (2000) describe how descriptive behavior of choice under uncertainty can violate this Bernoullian expectation rule. Nonlinear decision weights violate the invariance assumption, leading individuals to undervalue the reduction of a hazard probability relative to its complete elimination.

[28]In two-stage stochastic programs, the uncertainty resolution date refers to the period in which uncertainty is eliminated and information is revealed about the values of uncertain parameters.

adapt to many realized scenarios in a cost-effective manner. Adaptation costs can be viewed as the additional investments in the second stage that are required to adapt to the realized state of the world. In this sense, stochastic programming minimizes expected adaptation costs and, by extension, maximizes flexibility.

2.5 Capacity Planning Literature Review

Earlier sections in this chapter discussed frameworks for uncertainty analysis in energy modeling and highlighted the utility of sequential decision-making approaches with a focus on stochastic programming, which is the primary mathematical framework for projects in this dissertation. This section reviews the literature on applications of stochastic programming in energy modeling with a focus on electric sector capacity planning. The purposive sample of papers discussed here focuses on the central articles only and is not an exhaustive survey of all relevant papers. Emphasis is placed on work that relates to the research methods and application areas in Chapters 3 through 5—namely, models with a bottom-up technological representation, national scale, and expected-value decision criterion.[29]

2.5.1 Overview

Although previous sections underscored the limited formal inclusion of uncertainty in long-term energy models, this dearth is especially prominent for research using stochastic programming approaches with large-scale, bottom-up models. When employed, stochastic programming is most frequently implemented in aggregate global models for simple, *ad-hoc* modeling exercises with prominent examples including DICE (Nordhaus, 1994), MERGE (Manne and Richels, 1993), and CETA (Peck and Teisberg, 1993). The global scope and macroeconomic focus of these studies often come at the expense of technical detail about the energy system. The exclusion of such detail makes it difficult to extract insights about capacity planning and dispatch, specific technological uncertainties, or R&D strategy. This dissertation addresses

[29] Literature reviews related to R&D portfolio management (Chapter 6) and the impacts of fat-tailed uncertainty on climate policy (Chapter 7) are found in their respective chapters.

these concerns by developing an electric sector capacity planning model, which uses a stochastic programming approach with a high degree of technological detail.

Capacity planning models have a long history in the electric power sector, but generation expansion planning has only included uncertainty explicitly in recent decades. Long-range capacity planning for meeting a least-cost objective subject to operational constraints began in the 1950s alongside developments in mathematical programming (Massé and Gibrat, 1957). Even this early work suggests that effective capacity planning should treat uncertainty explicitly. Dapkus and Bowe (1984) present one of the first applications of stochastic programming to capacity planning in the electric power sector. The use of stochastic programming in energy models more generally, including long-term capacity planning and short-term hydro unit commitment decisions, is summarized in Wallace and Fleten (2003).

Tables 2.1 and 2.2 summarize the selected literature that will be discussed in subsequent sections. These tables give a sense of the scope of stochastic programming research related to electric sector capacity planning with technologically detailed models. The tables include relevant details about the studies, including the model, scope, aggregation, and uncertainties. The next section reviews stochastic capacity planning models with cost-minimization structures, and the subsequent section discusses optimal-growth models.

Table 2.1: Selected literature on stochastic programming models for electric sector capacity planning.

Paper	Model	Scope	Time Frame	OF
Birge and Rosa (1996)	Based on Global 2100	US	1990–2100	OG
Bistline and Weyant (2013)	Stochastic MARKAL	US	2000–2050	CM
Bosetti and Tavoni (2009)	WITCH	Global (12 regions)	2000–2100	OG
Bosetti et al. (2009)	WITCH	Global (12 regions)	2000–2100	OG
De Cian and Tavoni (2012)	WITCH	Global (12 regions)	2000–2100	OG
Durand-Lasserve, Pierru, and Smeers (2010)	Based on MERGE	Global	2005–2100	OG
Heinrich et al. (2007)	Stochastic MARKAL	South Africa	2002–2021	CM
Hu and Hobbs (2010)	Stochastic MARKAL	US	2000–2050	CM
Kanudia and Loulou (1999)	Stochastic MARKAL	Québec	1993–2037	CM
Kanudia and Shukla (1998)	Stochastic MARKAL	India	1993–2038	CM
Keppo and van der Zwaan (2012)	TIAM-ECN	Global (15 regions)	2010–2100	CM
Krukanont and Tezuka (2007)	N/A	Japan	2002–2014	CM
Labriet, Loulou, and Kanudia (2010)	ETSAP-TIAM	Global (15 regions)	2000–2100	CM
Loulou, Labriet, and Kanudia (2009)	ETSAP-TIAM	Global (15 regions)	2005–2100	CM
Manne and Richels (1993)	Global 2100	Global (5 regions)	1990–2050	OG
Peck and Teisberg (1993)	CETA	Global	2000–2200	OG
Usher and Strachan (2012)	Stochastic MARKAL	UK	2000–2050	CM

Notes: OF = objective function; CM = cost minimization; OG = optimal growth

Table 2.2: Uncertainties considered in the selected literature from Table 2.1.

Paper	Uncertain Parameter(s)	Number	Scenarios
Birge and Rosa (1996)	Returns on new technology investments, resource extraction, energy efficiency	1	4
Bistline and Weyant (2013)	Annual CO_2 caps, CCS availability, nuclear availability, CO_2 storage, natural gas prices	5	9
Bosetti and Tavoni (2009)	R&D effectiveness	1	3
Bosetti et al. (2009)	CO_2 concentration targets	1	3
De Cian and Tavoni (2012)	CO_2 taxes, capital costs	2	4
Durand-Lasserve, Pierru, and Smeers (2010)	Annual CO_2 caps	1	2
Heinrich et al. (2007)	End-use demand	1	3
Hu and Hobbs (2010)	Annual emissions caps (NO_x, SO_2, CO_2), end-use demand, natural gas prices	3	6
Kanudia and Loulou (1999)	Cumulative abatement targets, demand	2	4
Kanudia and Shukla (1998)	CO_2 taxes, end-use demand	2	9
Keppo and van der Zwaan (2012)	Radiative forcing targets, CO_2 storage	2	6
Krukanont and Tezuka (2007)	CO_2 taxes, demand (transport, electricity, heating), plant operating availability	6	6.2×10^{14}
Labriet, Loulou, and Kanudia (2010)	Climate sensitivity, economic growth	2	8
Loulou, Labriet, and Kanudia (2009)	Climate sensitivity	1	3
Manne and Richels (1993)	Annual CO_2 caps	1	3
Peck and Teisberg (1993)	Climate sensitivity, damage function	3	4
Usher and Strachan (2012)	Fuel prices, biomass import availability	2	4

Notes: Number = total number of uncertainties considered in the analysis; Scenarios = maximum number of scenarios (i.e., states of the world) considered simultaneously

2.5.2 Cost-Minimization Models

Cost-minimization models find feasible solutions that minimize the objective function of discounted energy system costs. Compared with their optimal-growth counterparts, cost-minimization models tend to have more detailed representations of energy technologies but do not capture macroeconomic interactions in as much detail.

The stochastic variant of the MARKet ALlocation (MARKAL) model is the most popular platform for the technology-detailed, policy-oriented studies in the literature.[30] Kanudia and Loulou (1998) investigate how the uncertainties of economic growth and mitigation can affect energy system planning for Québec.[31] A similar study by Kanudia and Shukla (1998) incorporates the same uncertainties for the Indian energy system and adds elastic demand. Heinrich et al. (2007) assess the impact of demand uncertainty on near-term decisions in a South African context using multiple objectives. Hu and Hobbs (2010) use stochastic MARKAL to calculate the EVPI, value of the stochastic solution (VSS), and value of policy coordination given uncertainties about multi-pollutant regulations in the United States (US), resource costs, and electricity demand. Usher and Strachan (2012) discuss near-term hedging strategies for long-term decarbonization pathways in the UK under uncertainties about fuel prices and biomass import availability. Bistline and Weyant (2013) demonstrate the utility of the stochastic programming framework and accompanying metrics using technological and policy-related uncertainties in the US electric sector as motivating examples. This paper draws attention to the limitations of stochastic MARKAL and stresses the need for new tools to better exploit the full range of benefits the stochastic programming approach can provide. Recent papers have used the stochastic TIMES model, which is an updated version of MARKAL, to examine uncertainties about the climate sensitivity parameter and economic development in multi-region, global settings (Labriet, Loulou, and Kanudia, 2010; Loulou, Labriet, and Kanudia, 2009).

The limitations of stochastic MARKAL are similar to other models employing stochastic programming in an energy policy context. Regardless of database size,

[30] An overview of the model can be found in Hu and Hobbs (2010).

[31] Loulou and Kanudia (1999) examine the same Québec system and compare results with a modified MARKAL model that uses a minimax-regret criterion.

stochastic MARKAL restricts the number of stages to two and the number of simultaneous scenarios to nine.[32] The latter constraint considerably limits the number of uncertainties that can be considered simultaneously and their degree of detail. The limited types of parameters that can be treated as random variables in MARKAL is also restrictive. The formulation only allows parameters like environmental bounds and demand to be treated stochastically (Kanudia and Loulou, 1999), which eliminates from consideration classes of problems with uncertain objective function coefficients (e.g., capital costs). Additionally, with the exception of Kanudia and Shukla (1998), the stochastic MARKAL studies do not include price-responsive demand, which can be limiting when examining the capacity planning problem under an uncertain climate policy.

Outside of the stochastic MARKAL framework, Krukanont and Tezuka (2007) analyze near-term investments and the EVPI in Japan under four uncertain policy regimes with uncertainties about CO_2 taxes, demand, and plant operating availability. Although the analysis considers more scenarios than other studies, the model uses an extremely short time horizon of 12 years, which is far shorter than the operating lifetimes of energy sector assets. Thus, the Krukanont and Tezuka (2007) study offers limited actionable insights about capacity planning decisions under uncertainty.

Finally, Keppo and van der Zwaan (2012) use the TIAM-ECN model to examine the impact of climate policy and CO_2 storage potential. The results suggest that, if a stringent climate policy is included, this possible scenario dominates the near-term strategy and that the climate policy uncertainty plays a more important role in mitigation timing than storage. Like the MARKAL studies, Keppo and van der Zwaan (2012) only account for a limited range of possible uncertainties and do not focus on capacity deployment decisions and the associated policy implications.

[32] Although the TIMES model allows for a maximum of 64 states of the world, the implementation is still based on directly solving the deterministic equivalent of the problem (Loulou and Lehtila, 2012), which severely limits the degree of model detail.

2.5.3 Optimal-Growth Models

In addition to cost-minimization models, some papers in the literature examine the impact of uncertainty on capacity planning using optimal economic growth models with utility-maximization objectives.

Manne and Richels (1993) are among the first to quantify the impact of uncertainty on energy decisions in the presence of climate change using a sequential decision-making approach. Using the Global 2100 model and an uncertain climate policy, their results show that the optimal near-term CO_2 emissions path lies between the extreme cases and that abatement levels are sensitive to the quality and timing of climate science research. Their results suggest that the value of perfect (or improved) information for reducing scientific uncertainty may be quite high and may be upward of $100 billion for the US. Like the cost-minimization studies, Manne and Richels (1993) only consider a very limited number of uncertainties and metrics for evaluation and instead focus on emissions trajectories.

Peck and Teisberg (1993) use the Carbon Emissions Trajectory Assessment (CETA) model to investigate uncertainty about the climate sensitivity parameter and damage function parameters. They find that information about these uncertainties has a large value relative to existing research budgets and that resolving uncertainty about impacts is almost as critical as learning about the climate sensitivity parameter. The benefits of resolving uncertainty early are considerably larger when suboptimal abatement is undertaken in the near term, assuming that the climate policy could be adjusted once more information is available. Again, the study's restricted concentration on climate-related uncertainty comes at the expense of insights related to technology-specific deployment.

Birge (1995) considers how uncertainty on investment returns for energy technologies affects economic output, consumption, and emissions under annual CO_2 restrictions. Using the VSS metric with a modified version of the Global 2100 model, the author finds that the optimal hedging strategy (i.e., instead of an expected-value strategy) increases economic output by approximately 1.4 percent annually and also recommends higher optimal CO_2 taxes in early periods. The results also indicate that an international market for CO_2 rights substantially lowers the VSS. Although the paper exemplifies how the VSS metric can be used for determining the impact of suboptimal planning on macroeconomic variables, Birge (1995) considers a very limited number of uncertainties and does not analyze how uncertainty affects technology-specific capacity installation and production decisions.

Bosetti and Tavoni (2009) analyze how innovation uncertainty may change climate policy recommendations using a stochastic variant of the World Induced Technical Change Hybrid (WITCH) model with a no-carbon backstop technology whose cost is a function of R&D spending. Results from an analytical model and WITCH conclude that accounting for uncertainty in R&D effectiveness decreases climate policy costs and increases R&D investments. However, the paper only considers a single cost-related uncertainty for a stylized technology under a fixed and certain climate stabilization policy (with an atmospheric CO_2 concentration target).

Bosetti et al. (2009) use WITCH to investigate the cost of uncertainty for global stabilization targets and quantify the economic costs associated with delayed abatement. The paper suggests that short-term inaction is the leading determinant of welfare losses and increased compliance costs for stringent policies. These results indicate that a moderate near-term policy would be an effective hedging strategy until new information about the long-term severity of climate change arrives, which mirrors other conclusions in the literature (Yohe, Andronova, and Schlesinger, 2004). The precautionary abatement under the hedging strategy is driven primarily by the stringent target of 450 ppmv (CO_2 only), which is explained by the convexity of marginal abatement costs. Energy efficiency investments are shown to be optimal hedges. Overall, the paper's focus on the macroeconomic impacts of suboptimal strategies crowds out investigations of technology deployment decisions and does not account for the effect of other simultaneous uncertainties on utility-scale decisions.

De Cian and Tavoni (2012) also employ the stochastic programming variant of WITCH with uncertainties associated with CO_2 taxes and capital costs for low-carbon technologies. Like the results in Chapter 4 of this dissertation, uncertainty about climate policy does not materially impact the first-stage abatement level or generation but primarily affects the portfolio of new capacity additions. The paper explores how different levels of uncertainty influence low-carbon capacity investments through mean-preserving spreads, which indicates that hedging in nuclear and renewables increases in CO_2 price uncertainty (and CCS investments decrease). While the paper gives a more thorough portrait of how climate policy uncertainty influences investments in individual technologies, the work ultimately only considers a limited number of potential policies, uncertainties, and metrics to evaluate the importance of uncertainty in decision-making.

Durand-Lasserve, Pierru, and Smeers (2010) illustrate how uncertainties about abatement targets (incorporated as annual CO_2 emissions caps) may impact near-term technology deployment decisions and CO_2 prices using a modified version of the MERGE model. The results show how this uncertainty can impact near-term capacity decisions and energy-sector prices on regional and global scales. Ultimately, the focus of the paper is on the relationship between global policy uncertainty, CO_2 prices, and emissions trajectories and not on capacity deployment decisions.

2.5.4 Discussion and Contributions

A few common conclusions emerge from comparing the studies discussed in the previous sections and listed in Table 2.1. First, results in the literature suggest that sequential decision-making approaches offer decision-relevant insights, which are not available through scenario analysis or Monte Carlo analysis. Stochastic programming models of capacity planning recommend strategies that differ from deterministic models where random variables are replaced by their expected values. Such results imply that the explicit inclusion of uncertainty is important for decisions, which demonstrates the utility of sequential decision-making frameworks.[33] Second, many studies point to the importance of climate policy uncertainties. Keppo and van der Zwaan (2012) conclude that uncertain climate targets dominate uncertainties about CO_2

[33]For instance, Usher and Strachan (2012) show that the stochastic hedging strategy is different from any deterministic (wait-and-see) solution and is structurally dissimilar from the average of the scenarios. Similar results are found in studies like Birge and Rosa (1996); Kanudia and Shukla (1998); Durand-Lasserve, Pierru, and Smeers (2010).

storage.[34] When climate policy uncertainty is included, it seems to impact the portfolio of deployed technologies more than near-term abatement levels (De Cian and Tavoni, 2012).

Another common thread in the literature is the low dimensionality of previous studies. As shown in Table 2.2, models typically consider a very limited number of uncertainties (with an average of 2.1 per study) and total scenarios (with an average of 4.7 per study).[35] The large number of decision variables and associated computational burdens have limited previous analyses to simple scenario trees, which prevent more than a couple uncertainties from being investigated simultaneously. The small number of scenarios capable of simultaneous representation in models like stochastic MARKAL restricts investigation of interactions between uncertainties.

The limited number of scenarios capable of consideration is a consequence of the solving approach. For stochastic linear programs with discrete distributions, the most common approach is to represent the problem as an equivalent deterministic linear program and then to solve directly, which is computationally costly for problems with many possible realizations.[36] Two-stage stochastic linear programs can take advantage of their special block structures through a variety of decomposition procedures (Birge and Louveaux, 2011). When the number of possible realizations of random parameters is particularly large, approximate solutions can be found through Monte Carlo sampling with variance reduction techniques (Infanger, 1999).

Table 2.3 shows that many studies in the literature do not take advantage of metrics for assessing the relative importance of uncertainties. As described in Chapter 3.3.3, metrics like the VSS and EVPI have important implications for decision-makers and modelers. However, the values in Table 2.3 suggest that only 29 percent of studies in the sample provide calculations for the VSS, 53 percent for the EVPI, and 6 percent for the value of control (VOC). Chapter 3.3.3 offers detailed discussions and mathematical definitions of these metrics.

[34] Additionally, when the potential for very stringent policy is included, Keppo and van der Zwaan (2012) conclude that it dominates the hedging strategy.

[35] The sample mean excludes the Krukanont and Tezuka (2007) paper, which considers notably more states of the world compared with other studies.

[36] This approach enumerates all variables and equations into single, large optimization problem.

Another shortcoming is the use of *ad-hoc* probability distributions instead of values based on rigorous modeling efforts. Many studies assume a uniform distribution over states of the world (i.e., invoking the Laplace criterion) to express an uninformative prior (Bosetti et al., 2009; Heinrich et al., 2007).

Table 2.3: Metrics used in the selected literature from Table 2.1.

Paper	**VSS**	**EVPI**	**VOC**
Birge and Rosa (1996)	•	•	
Bistline and Weyant (2013)	•	•	•
Bosetti and Tavoni (2009)			
Bosetti et al. (2009)			
De Cian and Tavoni (2012)			
Durand-Lasserve, Pierru, and Smeers (2010)			
Heinrich et al. (2007)	•		
Hu and Hobbs (2010)	•	•	
Kanudia and Loulou (1999)	•	•	
Kanudia and Shukla (1998)			
Keppo and van der Zwaan (2012)			
Krukanont and Tezuka (2007)		•	
Labriet, Loulou, and Kanudia (2010)		•	
Loulou, Labriet, and Kanudia (2009)			
Manne and Richels (1993)		•	
Peck and Teisberg (1993)		•	
Usher and Strachan (2012)		•	

Notes: VSS = value of the stochastic solution; EVPI = expected value of perfect information; VOC = value of control

Though some studies have made uncertainty analysis a central focus of capacity planning research (as illustrated in literature review above), the research in this dissertation is the first of its kind to investigate many uncertainties simultaneously by bridging state-of-the-art operations research techniques and a large-scale energy model. This research uses the DECIS system (Infanger, 1999), which is designed to use powerful decomposition techniques to solve stochastic programs with many scenarios. This modeling choice allows the research in Chapters 3 through 5 to incorporate a range of uncertainties with many thousands of scenarios.

The work here also proposes and applies a range of complementary metrics for quantifying the importance of uncertainty, which can indicate the value of reducing uncertainty and of using stochastic hedging approaches. These metrics provide a quantitative means of evaluating the significance of using sequential decision-making approaches for energy modeling in an environment of transition, where near-term decisions by economic agents seem to depend strongly on many uncertain factors.

Another contribution of this work is to offer rigorous quantifications of distributions for technological, economic, and policy-related uncertainties instead of using *ad-hoc* probabilities. This characteristic exhibits a higher degree of fidelity to the real-world investment problems of utilities and generators.

Additionally, previous studies often present numerical results from data-driven models but provide limited analysis and intuition for the underlying dynamics behind these outputs. In contrast, the work here analyzes results in greater depth in terms of the optionality of investments, which leads to insights about uncertainty, learning, and irreversibility in electric power sector. This research is also the first to incorporate parameters and uncertainties related to shale gas within a modeling framework that considers such risks explicitly. This feature provides unique and policy-relevant insights into the potential role of unconventional natural gas in the US electric power sector.

Other modeling contributions, which are discussed in greater depth throughout Chapter 3, include construction lead times, price-responsive demand, and the possibility of construction cost inflation.

2.6 Summary

This chapter described how more sophisticated treatments of uncertainty and the inclusion of sequential decision-making, especially through stochastic programming, can give insights into many dilemmas faced by decision-makers in energy and environmental domains. The discussion of stochastic programming illustrated the value of hedging strategies, which protect against a variety of risks and account for costs of midcourse corrections. The literature review discussed research gaps in applications of stochastic programming for electric sector capacity planning and highlighted contributions of the research in this thesis. Many of these contributions are presented in the context of model formulation and construction in the next chapter.

Chapter 3

Model of Capacity Planning under Uncertainty

Capacity planning in the electric power sector is well-suited to the stochastic programming paradigm, where strategies adjust over time as new information becomes available about technologies, resources, and polices. Decisions about capacity expansion and operation take place against long and highly uncertain planning horizons. Uncertainties about developments in the system environment impact the cost-effectiveness of planning decisions, particularly for utilities whose long-lived and essentially irreversible capital investments are designed to last many decades. The long lead times and lifetimes of energy assets mean that the environments in which power plants come online and operate may be very different from the ones in which they are planned. Hence, suboptimal near-term decisions that do not account for a range of potential fuel prices and environmental policies, for instance, can cost ratepayers, investors, and taxpayers and have important environmental implications.

Planning in the United States (US) electric power industry has been shrouded in substantial uncertainty in recent decades, and the simultaneous challenges with which the sector must grapple are only expected to increase in the future. Progressively stringent environmental policies, especially related to climate change, may require emission controls alongside early retirements and fuel switching. Complying with federal and state regulations must happen while utilities concurrently struggle with an aging fleet of generators and abrupt changes in the economics of fossil fuels due to the expansion of unconventional natural gas development. These factors make it even more important for decision-makers to develop strategies that hedge against a variety of possible futures and that explicitly consider both the expected costs and robustness of proposed plans.

This chapter formulates a two-stage stochastic programming model of capacity planning for the US electric power sector. The motivations of this model are to examine how technological, economic, and policy-related uncertainties can impact near-term planning decisions and to quantify the value of explicitly incorporating uncertainty and flexibility in the decision-making process. In addition to providing a comprehensive framework for analyzing adaptive management strategies, this model is among the first to use a stochastic programming framework in a large-scale energy-economic model with a wide range of simultaneous uncertainties and many scenarios. The model is also the first to incorporate upstream emissions from shale gas production into an energy-economic model that can examine tradeoffs between uncertain life-cycle costs and environmental impacts of different technologies.

This chapter presents the stochastic capacity planning and dispatch model by first formulating the deterministic version and discussing related framing assumptions in Sections 3.1 and 3.2, respectively. Section 3.3 describes the two-stage stochastic programming model and the definitions for metrics to quantify the importance of uncertainty. Finally, Section 3.4 considers the characterization of uncertainties included in the analysis.

3.1 Deterministic Capacity Expansion Model

This section discusses the development of an intertemporal capacity planning optimization model of the US electric power sector that addresses the aforementioned research questions. The discrete-time model determines optimal capacity investment and production decisions for the aggregate US electric sector between 2010 and 2050 in five-year increments with three load segments per year.[1] The segments create a

[1] The base year is 2010, which is used for calibration. The first projection period is 2015.

piecewise approximation of the load duration curve and preserve total annual generation and peak-load characteristics. The model uses a partial equilibrium framework with exogenous prices for most fuels. Data for the model come from a variety of public sources, as shown in Table 3.1.

Table 3.1: Data sources for capacity planning model inputs.

Data	Source
Capital and O&M costs	EIA (DOE/EIA, 2010)
Existing capacity	Form EIA-860 (DOE/EIA, 2011c)
Availability and capacity factors	EPA National MARKAL Database 2010
Fuel prices	EIA Annual Energy Outlook (DOE/EIA, 2011a)
Load	Based on Form EIA-860 (DOE/EIA, 2011c)

The model assumes that capacity installation and electricity production decisions are coordinated among utilities and generators. In the core deterministic model, utilities determine the path of investment and capital stock that minimizes the sum of discounted energy system costs for all capacity blocks during all periods while satisfying power system constraints.[2] The mathematical description of the model in this section uses the following sets and corresponding index notation:

Sets and Indices

$t \in \mathcal{T}$ time periods in planning horizon

$i \in \mathcal{I}$ generation technology types

$j \in \mathcal{J}$ load segments (i.e., subperiods in the load duration curve)

$s \in \mathcal{S}$ steps in the piecewise demand curve

The decision variables and parameters in the objective function are:[3]

[2]The model uses a discount rate of five percent unless otherwise noted to represent the market rate of return on capital.

[3]Decision variables for new capacity investments are continuous. The model does not include lumpy investments (i.e., large, discrete investments that are typically restricted to fixed sizes), which would require a mixed-integer formulation. The linear-programming formulation also does not account for economies and diseconomies of scale, which can be important in plant sizing decisions.

Decision Variables

x_i^t new capacity investment of generation technology i decided at time t (GW)

y_{ij}^t dispatched capacity of type i during load segment j at time t (GW)

w_i^t installed capacity of type i available at time t (GW)

u_s^t reduced demand from step s in the demand curve at time t (GW)

Parameters

δ^t discount factor at time t

c_i^t capital cost for type i at time t (\$/kW)

Δ_i construction delay for type i (years)

f_i^t total dispatch costs for type i at time t (\$/kWh)

τ_j^t duration of segment j at time t (hours)

g_i^t maintenance costs for type i at time t (\$/kW), including grid integration

p_s^t economic cost of reduced demand from step s at time t (\$/kW)

Given these variables and parameters, the linear cost-minimizing objective function (expressed in million \$) for the deterministic capacity planning problem is:

$$\sum_t \delta^t \left(\sum_i c_i^t x_i^{t-\Delta_i} + \sum_i \sum_j f_i^t \tau_j^t y_{ij}^t + \sum_i g_i^t w_i^t + \sum_s p_s^t u_s^t \right) \tag{3.1}$$

Thus, the four primary constituents of total costs are capital costs, dispatch costs, maintenance costs, and costs associated with reduced demand.[4]

The model explicitly represents a broad range of electricity generation technologies, including various generations of nuclear power, solar and wind technologies, electricity from biomass, and multiple forms of fossil-based generating technologies with a variety of fuels and carbon capture options. Technological cost and performance characteristics are exogenous inputs to the model, since the capacity planning formulation does not incorporate endogenous technical change. The model uses a

[4]Dispatch costs for preexisting and newly constructed generators are the sum of the variable operation and maintenance costs, fuel costs, and pollutant taxes. For carbon dioxide (CO_2) transport and storage costs, a piecewise supply curve for CO_2 storage is incorporated into the model and calibrated using data from Dooley et al. (2004).

vintaging structure to ensure that technological assumptions for a given time apply only to new deployments in that period and to create more realistic capital turnover and retirement dynamics.

The model incorporates a wind cost supply curve with increasing cost in deployed capacity. This curve accounts for the variable quality of wind resources in different regions of the country, heterogeneity in siting costs and availability, and interregional (and intraregional) transmission capacity constraints.[5]

All model variants include the following constraints:

- Load balancing (market-clearing condition)

$$\tau_j^t \left(\sum_i y_{ij}^t - \zeta_j^t \right) = \tau_j^t \left(d_j^t - \sum_s u_s^t \right) (1 + \alpha^t) \qquad \forall t, j \tag{3.2}$$

  where ζ_j^t represents net international exports during load segment j at time t (GW), d_j^t is the reference demand level (GW), and α^t is a factor that represents both transmission losses and a reserve buffer. This constraint ensures that demand is met in each subperiod and assumes that economical, grid-scale storage is not available.

- Dynamics of capital addition, turnover, and retirement

$$w_i^t = w_i^{t-1} + x_i^{t-\Delta_i} - x_i^{t-\Delta_i-L_i} \qquad \forall t, i \tag{3.3}$$

  where L_i is the lifetime of type i.

- Production capacity bounds

$$y_{ij}^t \leq a_{ij}^t w_i^t \qquad \forall t, i, j \tag{3.4}$$

  where a_{ij}^t represents the product of the availability factor (i.e., ratio of the amount of time a generator can produce electricity in a given period to the

[5]The curve is based on outputs from the Electric Power Research Institute's US-REGEN model, which endogenously determines transmission builds using detailed regional wind resource data, an hourly dispatch model, and trade between regions through cross-border transmission (Niemeyer et al., 2012).

period's duration) and capacity factor (i.e., ratio of the actual output of a generator to the available output at its full nameplate capacity) for type i during segment j at time t. This constraint mathematically formalizes the notion that unit dispatch cannot exceed the available capital stock.

- Demand reduction costs

$$p_s^t = p_0^t \left(d_j^t\right)^{-\frac{1}{\varepsilon}} \left(d_j^t - \frac{s}{n}\bar{r} d_j^t\right)^{\frac{1}{\varepsilon}} \qquad \forall t, s \tag{3.5}$$

where $\bar{r}$ is the maximum demand reduction (as a percentage of the reference value), n is the total number of steps in the stepwise linear representation of the aggregate demand curve, and ε is the own-price elasticity of demand at the end-use level.[6]

- Investment constraints

$$x_i^t \leq \bar{x}_i^t \qquad \forall t, i \tag{3.6}$$

where $\bar{x}_i^t$ represents the upper limit on new capacity investment of technology i at time t. These upper bounds on expansion are based on current pipeline or other technological constraints and signify real-world frictions for new capacity installations. These expansion constraints take the form of annual limits on investment in specific technologies (e.g., carbon capture is assumed to be unavailable before 2020) but also of cumulative bounds for technologies like wind, which has resource and siting constraints.

- Non-negativity constraints

$$x_i^t, y_{ij}^t, w_i^t, u_s^t \geq 0 \qquad \forall t, i, j, s \tag{3.7}$$

Since the electric power sector is characterized by long-lived and expensive investments, many technical and economic factors can contribute to the retirement of its costly generating assets. Retirements occur in the model through three mechanisms.

[6]This stepwise linear formulation of price-sensitive demand represents only price-induced energy conservation and efficiency, since autonomous conservation is implicitly included in the baseline load growth forecast. For a more thorough explanation of this approach, please refer to Kanudia and Shukla (1998).

First, capacity may retire endogenously through economic drivers when maintenance costs for units exceed their anticipated economic benefits. Second, units that are online at the beginning of the time horizon are likely to be fully depreciated before the end. These exogenous lifetime constraints for residual capacity are incorporated through an upper bound on the percentage of units of a particular type that are online in a given period. Finally, the third mechanism for retirements occurs when new capacity reaches its operating lifetime during the time horizon of the model run, which also represents an exogenous constraint based on unit lifetimes.

Optional constraints for model runs include climate policy constraints (cap and trade, carbon tax, or cumulative emissions cap), federal renewable portfolio standards, target wind penetration, constraints on investments for limited technology portfolio runs, and constraints to fix decision variables based on a reference run.

3.2 Assumptions

3.2.1 Utility Perspective

The capacity planning problem is framed from the perspective of utilities and generators in the aggregate US. The model assumes that capacity installation and electricity production decisions are centrally coordinated among all utilities and generators. This frame is akin to the representative agent assumption in many integrated assessments models, which posits a single decision-maker who has access to all of model's required input data and the complete authority to implement the model's recommended decision strategy.[7] Agents (i.e., utilities and generators) are endowed with identical beliefs about uncertain parameters in the capacity planning problem. The aggregate national focus is common among capacity planning models, which frequently adopt nonspatial frameworks where decisions about plant locations and investments in transmission equipment are excluded from the model.

[7] Although this assumption is partially true for production decisions on a regional level (i.e., in deregulated electricity markets where unit commitment decisions are determined by an Independent System Operator), it is not true on a national level for production decisions or for capacity planning decisions. In reality, capacity planning and dispatch decisions are decentralized and made by many heterogenous market participants.

This coordinated formulation provides a benchmark for system efficiency, which is valuable from a normative perspective even if it is unlikely and incomplete from a descriptive one. It serves as a benchmark for comparison much like competitive markets in welfare economics. By identifying the optimal cooperative outcome, the optimization model provides targets against which actual outcomes can be measured assuming that estimations of uncertainties are not systematically biased.[8] The coordinated, single agent model used here greatly simplifies the formulation of the optimization problem, which precludes the need for significant informational assumptions about utilities and generators.

As suggested by Equation 3.1, the objective function represents the sum of discounted systems costs and assumes risk neutrality on the part of the aggregate decision-maker. Examining aggregate system behavior from the perspective of a central utility planner makes the assumption of risk neutrality more plausible, because potential losses only represent a small fraction of overall wealth. The rationale is similar to the default assumption of risk neutrality for government projects described in the Office of Management and Budget's Circular A-4, which states that risk neutrality should be used unless reasonable grounds exist for alternate assumptions of risk aversion (OMB, 2003). This recommendation is especially relevant for the social optimizer perspective of the research and development (R&D) portfolio management work in Chapter 6.

Since optimal strategies depend strongly on the definition of the decision-makers and their framing of the decision problem, the utility perspective of the capacity expansion model has important implications for the formulation of the optimization problem and characterization of uncertainty. The objective of utilities and generators is to minimize cost while meet a variety of technical and economic constraints, which differs from social optimizers whose goal is to maximize the sum of producer and consumer surplus.[9] Thus, utilities do not account for the social cost of carbon in

[8]Stochastic dynamic programming research (Botterud, Ilic, and Wangensteen, 2005) suggests that decentralized decision-making for investments in generating capacity tends to underinvest in baseload capacity compared with centralized decision-makers.

[9]Transportation-related technologies are not included in this formulation. The utility frame for the capacity planning problem only includes transportation indirectly through the costs of serving

their objective function but only the portion of the externality that is internalized through policy. This distinction has significant implications for defining the climate policy uncertainty, including the selection of the state space under consideration and the probabilities attached to outcomes. Utilities consider climate policies themselves to be uncertain, whereas social optimizers consider the social cost of carbon to be uncertain. A limitation of other stochastic programming models of electricity capacity planning, as discussed in Chapter 2.5, is that they assume equivalence between the social planner's problem and the utilities' problem. Here, it is not assumed that the optimal climate policy will always be adopted to balance the marginal benefits and costs of mitigation.

3.2.2 Learning

There are a few types of learning, as described by Kolstad (1996) and Kann and Weyant (2000):

- *Autonomous (passive) learning* uses the passage of time to reduce uncertainty through simple observation. The exogenous arrival of information may occur at one time (as in this work, Manne and Richels (1993), and Nordhaus (1994)) or gradually over time (as in Kolstad (1996)).

- *Active learning* (i.e., learning from experience) uses observations on the states of a system to gain information about uncertainty. This type of learning monitors the impacts of decisions on variables of interest to gain knowledge about uncertain parameters. For many issues in energy and environmental policy, active learning experiments are challenging due to detection difficulties, time lags, and irreversibilities.[10]

increased demand from electrified vehicles, which is considered in the sensitivity analysis in Chapter 4.2.6. The results of Usher and Strachan (2012) indicate that near-term hedging decision related to transportation are very similar to the deterministic, expected-value strategies.

[10]For instance, Kelly and Kolstad (1999) investigate how perturbing greenhouse gas emissions can convey information about climate-related uncertainties. They find a "tradeoff between the expected benefits of controlling emissions and the resolution of uncertainty," since emissions controls lead to slower learning.

- *Purchased learning* occurs when knowledge can be obtained through mechanisms like research. The optimal expenditure level on purchased learning depends on the costs, benefits, and uncertainties of information gathering in a specific decision context.

Stochastic programming frameworks typically represent autonomous learning. The information structure for such two- or three-stage models contains priors on states of the world, a vector of probabilities associated with receiving specific messages, and *ex-post* probabilities for states of the world conditional on specific messages. Although *ex-ante* probability distributions are known, the *ex-post* marginal distribution is known only after observations are made at the onset of the second stage.

Thus, the two-stage capacity planning model formulated in this chapter and the corresponding results in Chapters 4 and 5 involve autonomous learning with perfect information that arrives during a single period instead of gradually over time. Chapter 6 considers the modeling and policy implications of purchased learning in the context of energy technology R&D portfolio management.

3.3 Two-Stage Stochastic Programming Model with Recourse

The linear programming model discussed above computes the optimal investment and operational strategies for the deterministic capacity expansion problem. Under perfect information, this solution provides a lower bound on discounted costs given a particular scenario. However, due to the difficulties associated with predicting the outcomes of uncertainties introduced in Section 3.4, it is unrealistic to assume that a strategy that is optimized for a given scenario will be optimal under a range of realized states of the world. Disregarding inherently random characteristics may limit the usefulness of solutions designed using deterministic approaches.

3.3.1 Formulation

As discussed in Chapter 2, stochastic programming techniques can be used to compute optimal hedging strategies in problems with uncertain data and to provide contingency plans that adapt to realizations of random variables (Beale, 1955; Dantzig, 1955). These solutions perform reasonably well under a variety of plausible scenarios. The basic two-stage stochastic program with recourse (Birge and Louveaux, 2011; Infanger, 1994) in a cost-minimization setting can be formulated as:

$$\begin{array}{lrlrl} \min & z = & c^T x + & \mathbb{E}_\omega f^\omega y^\omega & \\ \text{s.t.} & & Ax & & = b \\ & & -B^\omega x + & D^\omega y^\omega & = d^\omega \\ & & x, & y^\omega & \geq 0,\ \omega \in \Omega \end{array}$$

Ω	set of all outcome paths
$\omega \in \Omega$	state of the world
x	vector of first-stage decisions
y	vector of second-stage (recourse) decisions

Here, all values corresponding to objective function coefficients (i.e., the c vector) and first-stage constraints (i.e., the A matrix and b vector) are known with certainty. The second-stage objective coefficients (i.e., the f^ω vector) and parameters in the constraints (i.e., the B^ω and D^ω matrices and d^ω vector) are unknown when utilities make first-stage decisions and are characterized only by discrete probability distributions over potential outcomes.[11] The second-stage parameters are treated as random variables with outcomes denoted by ω with an associated probability $p(\omega)$. Every random element depends jointly on these scenarios or states of the world.

This framework requires the specification of a full set of random variable realizations and the corresponding probabilities of occurrence. The set Ω represents

[11] B is often called the transition matrix, and D is known as the technology or recourse matrix. The superscript on the second-stage vector y^ω makes explicit that the choice of recourse decision variables depends on the realization of scenario ω.

a multiplicity of potential futures and contains a full panoply of outcomes from the uncertainties in Section 3.4. Realizations of random variables, designated by the symbol ω, are called scenarios or states of the world.[12] Outcomes of Ω can be grouped into subsets known as events with a σ-algebra $\mathcal{F} \subseteq \Omega$. The probability measure $P \colon \mathcal{F} \to [0,1]$ on the space $(\Omega, \mathcal{F})$ is a real-valued function that satisfies the countable additivity property and $P(\Omega) = 1$. Thus, the mathematical triplet $(\Omega, \mathcal{F}, P)$ represents the probability space for the model.

The objective is then to identify a point in the space of all admissible values of decision variables that corresponds to the extremum of the objective function. In this case, utilities and generators are minimizing the expected sum of discounted energy system costs subject to the many techno-economic constraints of the energy system. The optimal first-stage strategy minimizes expected costs, which consist of first-stage costs $c^T x$ and the expected costs from recourse decisions $\mathbb{E}_\omega f^\omega y^\omega$ (i.e., the penalty for correcting first-stage decisions after uncertainties are resolved).

3.3.2 Solution Approaches

In decision contexts under uncertainty, two important questions for decision-makers and modelers are:

- How much should decision-makers be willing to pay for information about uncertain quantities?
- What is the value of incorporating uncertainty explicitly in the decision-making process instead of using a deterministic approximation?

Stochastic programming settings provide convenient mathematical frameworks for defining and quantifying answers to these questions.

Before discussing these metrics used to evaluate the impacts of uncertainty, it is useful to distinguish between three approaches for solving decision problems under uncertainty (Infanger, 1994). The *wait-and-see (learn-then-act) approach* assumes that

[12] Again, although decision-makers cannot completely observe the path of uncertain parameter values (i.e., the random element ω) when making decisions, this framework assumes that the probability distributions over outcomes are known.

uncertainty is resolved and the outcome $\omega \in \Omega$ can be observed before selecting the optimal decision vector x (Mandansky, 1960). This solution corresponds to a scenario analysis problem (i.e., where uncertainty has been removed and the decision-maker solves for different values of ω) and suggests that perfect information is available before decisions are made. The problem can be formulated as:

$$
\begin{aligned}
z^{\omega} = \min\ & f(x,\omega) \\
\text{s.t. } & x \in C^{\omega} \subseteq \mathbb{R}^{n}
\end{aligned}
$$

with the wait-and-see solution expressed as $x^{\omega} \in \operatorname{argmin}\ \{f(x,\omega) \mid x \in C^{\omega}\}$. The expected cost with perfect information can be found by taking the expected value over all possible scenarios: $z_{ws} = \mathbb{E}\, z^{\omega} = \sum_{\omega \in \Omega} z^{\omega} p(\omega)$. The problem with this approach is that the solutions are not implementable, which means that outcomes are superoptimal and decisions cannot be regretted *ex post*. Although the assumption that learning will resolve uncertainty completely is unrealistic in many contexts, the wait-and-see approach serves as a valuable conceptual benchmark against which expected costs of other strategies can be measured and compared, which makes it the basis for value of information calculations.

Second, the *here-and-now approach* finds a solution x^* that hedges against all possible contingencies $\omega \in \Omega$ that may occur in the future.[13] The optimal stochastic solution addresses the decision problem where a decision-maker cannot completely resolve uncertainty before acting, which closely resembles many decision contexts. This decision is made before observing the outcome from Ω and solves the problem:

$$
\begin{aligned}
z^{*} = \min\ & \mathbb{E}_{\omega} f(x,\omega) \\
\text{s.t. } & x \in C^{\omega} = \bigcap_{\omega \in \Omega} C^{\omega}
\end{aligned}
$$

where the stochastic solution is expressed as $x^* \in \operatorname{argmin}\ \{\mathbb{E}_{\omega} f(x,\omega) \mid x \in \cap C^{\omega}\}$. The solution x must be feasible for all scenarios $\omega \in \Omega$. The expected cost of the

[13]This is also called the *act-then-learn approach* in problems without recourse actions and the *act-then-learn-then-act approach* when recourse actions are available in the second stage.

stochastic solution is $z^* = \min_x \mathbb{E}_\omega f(x,\omega)$. The difficulty with this approach is its computationally intensive nature.

Finally, the *expected-value approach* replaces the stochastic parameters by their expected values or another measure of central tendency like the median or mode.[14] This approach sidesteps uncertainty by using a single set of input parameters and solves the problem:

$$\hat{z}_d = \min f(x,\bar{\omega})$$
$$\text{s.t. } x \in C^{\bar{\omega}}$$

where the expected-value solution is $x_d \in \text{argmin}\,\{f(x,\bar{\omega}) \mid x \in C^{\bar{\omega}}\}$ and $\bar{\omega} = \mathbb{E}\omega = \sum_{\omega\in\Omega} \omega p(\omega)$. The expected cost of the expected-value solution is $z_d = \mathbb{E}_\omega f(x_d,\omega)$. Much like the wait-and-see approach, the expected cost of the expected-value solution can be found using a deterministic model. This approach can be beneficial due to its computational ease in formulating and solving a deterministic problem that approximates the actual decision problem. However, this simplification may exclude critical dynamics of the uncertain system and may perform poorly in expectation, as quantified through the value of the stochastic solution.

Figure 3.1 compares the expected costs of the expected-value (z_d), stochastic (z^*), perfect information (z_{ws}), and control (z_c) strategies.[15] The figure also illustrates the metrics discussed in Section 3.3.3.

[14] If the probability density function is roughly symmetric and unimodal, the choice of a measure of central tendency (i.e., mean, median, or mode) does not make a substantial difference in results. However, policy-relevant uncertainties in many system involve distributions with nonneglible skewness and kurtosis.

[15] Mandansky (1960) proves that $\hat{z}_d \leq z_{ws} \leq z^* \leq z_d$ using Jensen's inequality and the convexity of the objective function.

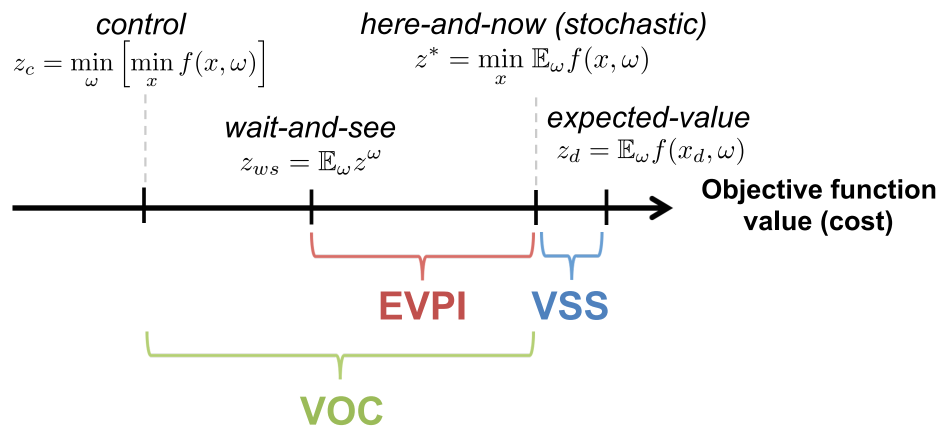

Figure 3.1: Number line comparing expected costs under different decision-making approaches. The spacing between values is illustrative.

The input parameters for the expected-value approach do not have to be the expectations of the random variables in a model. Instead, this reference or best-estimate scenario can consist of any single nominal value for all input parameters based on some central measure as defined by the decision-maker. Thus, the reference scenario approach can be interpreted as a heuristic strategy that uses approximations for parameter values. Many decisions are not made by explicitly accounting for uncertainty in sophisticated analytical models but are informed by a combination of experience, rules of thumb, *ad-hoc* heuristics, and luck. As described in the results in Chapter 4, the selection of this reference case can make a large difference in calculating the value of the stochastic solution.

There are many reasons why a decision-maker may choose to use a set of parameters that differs from expectations of random variables:

1. The decision-maker may not be aware of an uncertainty or may not consider the parameter to be uncertain and consequently does not include this possibility in its calculations. For instance, public acceptance uncertainties surrounding large-scale CO_2 storage is rarely incorporated into capacity planning analyses. With many simultaneous uncertainties with which to contend, decision-makers can be overwhelmed with the number of factors that must be taken into consideration. The result may be to adopt more lax tools for incorporating secondary

uncertainties into the planning process.

2. The decision-maker may be aware of an uncertainty but does not have the time or resources to characterize probability distributions and then to calculate the expectation explicitly. As a result, planning may use best guesses instead of expected values.

3. The decision-maker may be aware of an uncertainty and have the resources to quantify the expected value, but cognitive heuristics and biases may distort the distribution from the "true" distribution. For instance, optimism biases could underestimate planning and operational costs (Kahneman and Tversky, 1982), or overconfidence may exclude important tail events (Tversky and Kahneman, 1982). Perception and management of risk can be distorted by a range of biases, which consequently can impede a decision-maker's ability to quantify and prioritize risks accurately.

4. The decision-maker may be aware of an uncertainty and have an unbiased distribution over possible outcomes but may view that uncertainty as endogenous with (partially) controllable outcomes. For instance, a utility or generator may plan under the assumptions of no climate policy or a lenient one, believing that they can influence or indefinitely delay carbon-pricing policies or can achieve special exemptions.

3.3.3 Metrics

The importance of uncertainties is assessed through three metrics: the *expected value of perfect information* (EVPI), *value of the stochastic solution* (VSS), and *value of control* (VOC). Of the limited energy modeling research that uses a stochastic programming framework, few studies perform EVPI calculations and even fewer incorporate VSS metrics, as discussed in Chapter 2.5. No studies explicitly use the VOC metric or draw attention to its relation to R&D activities. The underutilization of these metrics neglects important opportunities for policy insight and model development afforded by stochastic programming.

This section defines these metrics mathematically and describes how modelers can exploit them to improve policy recommendations as well as model development and diagnostics, as demonstrated Chapter 4.

Expected Value of Perfect Information

First defined by Raiffa and Schlaifer (1961), the EVPI compares the expected costs of the stochastic and wait-and-see solutions and represents the expected change in the objective function value if perfectly accurate forecasts are available prior to first-stage decisions.[16] The EVPI has important implications for decision-makers in placing an upper bound on their willingness to pay for information-gathering activities.[17] This value can help to establish limits during the budget allocation process in support of research programs to improve knowledge of uncertainties (e.g., forecasting research). It quantifies the potential value added from determining which outcome will actually occur, which avoids losses from irreversible investments and opportunity costs from delay. The EVPI is mathematically defined as:

$$\begin{aligned}\text{EVPI} &\equiv z^* - z_{ws}\\ &= \min_x \mathbb{E}_\omega f(x,\omega) - \mathbb{E}_\omega\left[\min_x f(x,\omega)\right]\end{aligned}$$

The schism between the here-and-now and wait-and-see approaches arises from the fact that it is unlikely for a single solution to be both feasible and optimal for every scenario. In this sense, the EVPI can be interpreted as the cost of uncertainty incurred by a decision-maker when following an optimal hedging strategy. Since the wait-and-see approach always yields a better outcome (or, strictly speaking, no worse outcome) under any scenario compared with the here-and-now approach, the EVPI is always nonnegative. The sufficient condition to ensure a zero-valued EVPI is for the optimal stochastic solution to be independent of uncertain model parameters (Birge and Louveaux, 2011).

[16] The EVPI is also called the "value of clairvoyance" in decision analysis parlance (Howard, 1968).

[17] In many applications, information is neither complete nor perfectly accurate, and it may be impractical to assume that uncertainty can be completely eliminated. As a result, the expected value of imperfect information is less than the EVPI.

For the EVPI calculations, observations are assumed to be both pure and perfect (Matheson and Matheson, 2005). Pure observations implicitly assume that the act of gathering information and making observations leaves everything else in the influence diagram (i.e., generalized Bayesian network) unchanged. Reports from perfect observations assign probabilities of one to the realized outcomes of target nodes.

The existence of the EVPI is predicated on the notion that there may be explicit or implicit opportunity costs associated with delaying decisions. Decisions based on full information generally perform better than ones with incomplete information *ceteris paribus*. If a decision-maker can costlessly delay action until uncertainty is resolved, then inaction will be a preferred strategy. The decision about the extent to which resources should be committed now instead of in the future when information may be available depends on a host of factors, including the riskiness and opportunity costs associated with delay. If delay is risky, this strategy may not be sensible, which gives rise to the EVPI due to differences between the wait-and-see and here-and-now decisions during the first stage. Thus, the EVPI is a proxy for the opportunity cost of delaying a decision until more information is available.[18]

This metric has significant consequences for decision-makers. In addition to providing a bound on potential benefits from prediction and forecasting,[19] the EVPI can be interpreted as the expected regret of the here-and-now strategy.[20] Since it is the probability-weighted sum of regret in all possible states, the EVPI can be conceptualized as expected regret when regret is defined as the difference in payoffs between the here-and-now and wait-and-see (i.e., perfect information) strategies. Thus, the optimal stochastic solution—in addition to minimizing the expected value of the objective function—minimizes expected regret. This strategy is different from the minimax-regret approach, which minimizes the maximum regret over all states. When the EVPI for a specific uncertainty is nonzero, it indicates that no true "no

[18]Conrad (1980) demonstrates the equivalence of the value of information and the Arrow-Fisher-Hanemann-Henry option value.

[19]This allows decision-makers to ascertain whether the costs of gathering information to tighten distributions are justified by their potential benefits.

[20]Regret is the cost to the decision-maker for making a planning decision that is mismatched with the realized state of the world.

regrets" strategies exist across all decision variables for that uncertainty.[21]

Additionally, the EVPI captures the notion that, even though parameters may contribute to variations in the objective function, these uncertainties may be irrelevant to decisions. This feature allows stochastic programming models to identify policy-relevant parameters, which is one of the limitations to deterministic sensitivity analyses and uncertainty propagation approaches.

The EVPI also has implications for modelers. In contexts where the EVPI for a specific uncertainty is small but the VSS is large, it signals to analysts that resolving uncertainty is not as important as accounting for uncertainty explicitly in modeling efforts through sequential decision-making frameworks. These cases occur when suboptimal policies (e.g., those based on expectations of parameter values) lead to significant losses. A second point is that the act of constructing more detailed and elaborate models for uncertain quantities can lead to a greater understanding of the complex system itself and to reductions in uncertainty. The EVPI can be adapted as a sensitivity measure to estimate the maximum expected value of such forecast modeling and to suggest which uncertain features are worth developing in greater detail (Morgan and Henrion, 1990).

Despite the usefulness of this metric, there are a few important questions and complicating factors involved in calculating the EVPI. First, synergistic effects among contemporaneous uncertainties mean that the joint value of information may differ from the sum of individual ones, which makes such effects important areas of exploration. Second, influence arrows on decision diagrams (i.e., ones from decisions to uncertainties) make EVPI calculations slightly more difficult, as reports must be conditioned on each (discrete) decision alternative to avoid loops on the influence diagram. Decision structures in which decisions influence uncertainties are common in R&D management where, for instance, allocation decisions influence the probability of successful program outcomes. Third, it is often challenging to identify a decision-maker who would place a value on information or to pinpoint appropriate sources information for all uncertainties. In some cases, uncertainty may be irreducible, even

[21] A zero-valued EVPI is a necessary but not sufficient condition for a "no regrets" strategy. The strict definition is that $z_{ws}^{\omega} = z_{*}^{\omega}$ for all $\omega \in \Omega$.

though the EVPI is relatively large. Additionally, although one hopes that gathering more information about a problem will reduce uncertainty, research and exploration often reveal unforeseen complexities and temporarily can increase uncertainty.[22]

Finally, since the EVPI depends on prior information about an uncertain parameter (in addition to the alternatives available to a decision-maker and the decision-maker's objective function), the EVPI can be extremely sensitive to tail probabilities of prior distributions. If the well-established overconfidence effect narrows prior distributions, the EVPI is likely biased downward, which may lead decision-makers to substantially underestimate the utility of information gathering (Hammitt and Shlyakhter, 1999).

Value of the Stochastic Solution

The VSS compares the expected costs of the expected-value and stochastic solutions and describes the "value of knowing how little you know" (Morgan and Henrion, 1990).[23] It quantifies the expected difference in cost for a decision based on stochastic analysis and one that ignores uncertainty by opting to use a deterministic solution. The VSS can guide analysts in the process of model construction by highlighting which uncertainties are most important for inclusion and for more detailed probability elicitations. The VSS is defined by the equation:

$$\begin{aligned} \text{VSS} &\equiv z_d - z^* \\ &= \mathbb{E}_\omega f(x_d, \omega) - \min_x \mathbb{E}_\omega f(x, \omega) \end{aligned}$$

The explicit inclusion of uncertainty in planning efforts makes decision-makers at least no worse off in expectation, assuming that the additional costs of analysis and

[22]For example, Morgan and Keith (1995) asked climate experts to assess the probability that uncertainty about the value of the climate sensitivity parameter would grow by 25 percent or more after a 15-year research program with a $1 billion research budget annually. The average value from respondents was 19 percent. Chapter 7 discusses issues related to learning and uncertainty in the context of climate policy.

[23]The value of the stochastic solution is also called the "expected value of including uncertainty" (Morgan and Henrion, 1990). As discussed in greater depth in Section 3.3.2, the stochastic solution also may be compared to other reference scenario values that use different best estimates (e.g., other measures of central tendency like the median or mode) instead of the mean value.

interpretation are reasonably negligible. Uncertain prior distributions for parameters contain potentially useful information that otherwise would be lost if such values were assumed to be certain. Thus, like the EVPI, the VSS is always nonnegative.

The value of using a stochastic solution comes from a variety of features in the uncertain decision problem, including asymmetric loss functions, strongly asymmetric probability distributions for uncertain quantities, and dependence between random variables (Morgan and Henrion, 1990). The VSS is largest in contexts with asymmetrical payoffs for miscalculations and misforecasts. These situations with catastrophic loss functions occur when harms increase proportionately with error on one side of the optimal decision value but have a step-function-like error on the other side, which leads to a substantial and abrupt losses for even small deviations. For instance, the height of levees and seawalls exhibit highly asymmetrical payoffs during floods, tsunamis, or other events. If water levels are below the design height, losses are relatively small; however, if water overtops the crest of a levee or sea wall, damages are much more extensive, particularly if overtopping leads to a complete breach.[24]

The VSS has numerous implications for decision-makers. Since stochastic programs are more difficult to formulate and solve than deterministic models, the VSS can indicate whether approximations of optimal strategies (e.g., expected-value solutions) are nearly optimal or extremely suboptimal. Hence, decision-makers can use the VSS as a means to gauge the quality of an approximate solution and to quantify the value of incorporating uncertainty explicitly instead of assuming a certain value and then solving a deterministic decision problem.[25]

Additionally, the VSS can identify uncertainties that induce "anticipatory actions" (Labriet, Loulou, and Kanudia, 2010) and merit explicit treatment through a sequential framework. Random events do not generate anticipatory actions if the VSS is zero, which means that the hedging strategy is unresponsive to the uncertainty. A high VSS indicates the need for precautionary hedging and explicit treatment using

[24] Morgan and Henrion (1990) speculate that the pervasive use of linear and quadratic loss functions account for the widespread belief that explicitly incorporating uncertainty into decision support models does offer significant improvements over deterministic frameworks.

[25] The VSS can be interpreted as the value of what Keats calls "Negative Capability," which refers to the ability to tolerate uncertainty without "irritable reaching" for certainty (Keats, 1899).

a sequential approach, whereas other uncertainties can be treated in a simple sensitivity analysis. Thus, the question of how to detect events that induce anticipatory actions, as asked in Labriet, Loulou, and Kanudia (2010), is answered through the VSS metric.

The VSS also has important consequences for modelers. It quantifies the added value of using a stochastic, sequential decision-making model and guides analysts in the process of model construction by identifying the most important uncertainties for inclusion. Since the VSS quantifies the degree to which approximations of hedging strategies (e.g., expected-value solutions) are suboptimal, the use of this metric is especially valuable for utility resource planning, since it is common to approximate stochastic models by solving different deterministic models based on a range of parameter assumptions. Stochastic hedging strategies are then approximated in an *ad-hoc* manner by finding common elements of decision strategies for many states of the world (Jin et al., 2011). The VSS also can be used to prioritize efforts to quantify various uncertainties by offering general guidance for tasks like allocating resources across a range of modeling projects to assess distributions (e.g., building econometric models), selecting the extent and exhaustiveness of elicitation efforts, or determining computational tradeoffs when discretizing distributions.

This metric may be most relevant to modelers in contexts where it is difficult or impossible to obtain accurate information about the future (Birge, 1995). In these situations, adopting sequential decision-making can hedge against a range of potential futures. Deterministic formulations, especially the common approach of using linear programs, are prone to corner solutions (sometimes called "knife-edge" or "bang-bang" solutions) that provide little robustness. Sequential decision-making frameworks typically develop hedging solutions with a more diverse set of decision variables and with a greater degree of slackness to forestall costly adjustment penalties if realizations of uncertainties deviate from their expected values. This "practical validation advantage" (Birge, 1995) of stochastic programming means that decision-makers are more likely to view proposed strategies as being appropriately diversified, which makes such solutions less likely to be modified or rejected. This intuitive appeal is critical when the modeler and decision-maker are different people.

Initially, it appears contradictory that the VSS can be used as a means of determining whether a stochastic modeling framework is appropriate, since calculating the VSS seems to assume that a probabilistic model is already available. However, upper bounds for VSS values can be approximated using a core deterministic model to calculate whether a stochastic approach would be appropriate for a particular setting.[26] Birge (1982) shows how the EVPI and VSS are bounded by $z_d - z_{ws}$. The quantities needed to calculate this bound, the expected cost of the expected-value solution (z_d) and wait-and-see solution (z_{ws}), can be computed using standard deterministic models, provided that the modeler can fix first-stage decisions when calculating z_d and that probability distributions are available for uncertain parameters. This bound allows modelers to determine whether a specific application merits converting a deterministic model into one with a more sophisticated treatment of uncertainty. If the upper bound is substantial, more investigation may be necessary; however, if the bound is negligible, the expected-value strategy can be employed without a significant loss of fidelity.

Even if a stochastic model is available, the VSS is still useful for model development. The VSS for each uncertainty under consideration can be calculated individually to determine which random parameters have the largest VSS magnitudes, which means that they should be included in more computationally complex joint uncertainty model runs and should be given additional attention during the uncertainty quantification and model refinement processes.

The EVPI and VSS are different metrics that compare the expected value of the stochastic solution with another made without incorporating uncertainty. For the VSS, the other decision is made when uncertainty is disregarded, even though it still exists. For the EVPI, the other decision is made when uncertainty is removed by obtaining perfect information. Thus, the VSS can be viewed as the additional expected cost of pretending that uncertainty does not exist, whereas the EVPI is the expected cost of being uncertain (Morgan and Henrion, 1990). Seen differently, the VSS quantifies the value of incorporating uncertainty, and the EVPI "measures the reward for resolving uncertainty" (Birge and Rosa, 1996). The VSS and EVPI

[26] The risk premium may also provide an upper bound on the VSS (Morgan and Henrion, 1990).

metrics for an uncertainty can have very different magnitudes, as large values of one metric do not guarantee high values for the another.[27] Both metrics are expressed in units of the objective function.

Value of Control

The VOC is a useful metric for measuring the value of being able to control the outcome of an uncertain situation.[28] In some contexts, uncertain parameters can be considered decision variables if it is possible to exert some degree of control over its outcome by committing resources. The VOC represents the change in value for moving from an uncertain state (where the here-and-now approach is optimal) to a desired state without uncertainty. Assuming perfect control, the VOC is determined through the equation:

$$\mathrm{VOC} \equiv \min_x \mathbb{E}_\omega f(x,\omega) - \min_\omega \left[\min_x f(x,\omega) \right] \tag{3.8}$$

The VOC is useful for uncertain parameters that are controllable (either wholly or in part) through allocation decisions. These endogenous uncertainties manifest themselves, for instance, in the energy policy domain through technologies that have cost and performance characteristics that can be influenced through directed R&D efforts.

For the VOC calculation in Equation 3.8, control interventions are assumed to be both pure and perfect (Matheson and Matheson, 2005). Pure interventions only change probability assessments of the target node and leave the remainder of the influence diagram unchanged. Perfect interventions achieve its intended result completely, giving the decision-maker certain control over the outcome instead of merely shifting the probability distribution.

[27] Birge (1982) demonstrates how these metrics are "distinct, different values that measure different types of uncertainty" and proves that one metric can be zero while the other is nonzero.

[28] The VOC is known as the "value of wizardry" in decision analysis (Howard, 1971).

3.4 Uncertainties Considered in Analysis

This analysis represents eight uncertain model parameters as random variables, including the stringency of a federal climate policy, natural gas price path, coal price path, upstream methane (CH_4) emissions from shale gas, capital costs for nuclear and coal with carbon capture and storage (CCS), public acceptance of CO_2 storage, and natural gas combined cycle efficiency. These chosen uncertainties come from a variety of resources like informal interviews with experts in the utility industry and academia, existing peer-reviewed research, and integrated resource plans (IRPs) for utilities. This exploratory screening process identified the most important uncertainties to consider for inclusion. The subset of these uncertainties that was ultimately chosen for the final analysis was based on a series of sensitivity analyses, tornado diagrams, and approximations of the metrics in Section 3.3.3. Uncertainty is characterized for quantities in the model using approaches like statistical analyses, expert elicitations, and econometric modeling.

This section describes the motivations for including these uncertainties and how their finitely supported distributions are selected.

3.4.1 Climate Policy

Although global climate change is an urgent and significant problem, there are many sources of uncertainty that will determine the stringency of policy measures used to curb greenhouse emissions in the US and elsewhere (IPCC, 2007). The portfolio composition of generating assets for utilities over the coming decades will be influenced critically by firms' expectations about the timing, form, breadth, and stringency of future climate policy. Climate policy is one of the most important uncertainties currently facing the power sector due to the industry's heavy reliance on carbon-intensive generation, which creates significant exposure to climate policy risk. Although there are currently no federal regulations constraining greenhouse gas emissions, regulation and/or legislation is expected in the near future, and its characteristics are highly uncertain. Thus, it is appropriate and necessary for utilities and generators to consider strategies that reduce the exposure of ratepayers and shareholders (in some cases) to the cost risks associated with future policies.

The future form of a climate policy is unknown. Legislation could come in the form of a quantity instrument (e.g., quotas, commands, or targets) or price instrument (e.g., tax). Additionally, the potential for qualitative and quantitative restrictions in specific bills make it difficult to forecast the availability of diverse compliance mechanisms even if there were more certainty about the form of the legislative instrument itself.[29] If the legislative route stalls, regulation through the Environmental Protection Agency (EPA) is also possible after the US Supreme Court found that greenhouse gases are covered by the Clean Air Act and, through the April 2007 *Massachusetts v. EPA* decision, ruled that the EPA has the authority to regulate greenhouse gas emissions as pollutants.

It is important to note that these carbon taxes are fundamentally different from the social cost of carbon, as a variety of barriers may prevent the socially optimal policy from matching the one that is actually implemented. Distortions may arise due to many factors outside of the scope of this research, which can mean that the imposed tax does not reflect the actual benefits or costs of mitigation. Firms are indifferent toward the magnitude of external damages apart from the degree to which such externalities are internalized through policy.[30] Lobbying efforts by energy firms are one source of distortions, as some regulated entities may have financial incentives to push for more lenient legislation or for exemptions. For instance, lobbyists for fossil-fuel-related industries spent $500 million in 2009 and 2010 to defeat the Waxman-Markey bill, outspending proponents of the bill by a ten-to-one margin (Wagner and Zeckhauser, 2012).

The model does not consider other regulations like water, particulate matter, ozone, ash, or regional haze. The only greenhouse gases incorporated in the model are CO_2 and CH_4. It is assumed that the climate policy will include both of these greenhouse gases.

This analysis assumes that the policy will take the form of a price on equivalent

[29]Bills like Waxman-Markey (H.R. 2454) allow utilities to meet compliance obligations in part through certified offset credits.

[30]This framework does not model strategic interactions between firms and regulators, which is investigated by Tarui and Polasky (2005).

greenhouse gas emissions, which is uncertain until the second stage. After this period, the CO_2 price is assumed to escalate in real terms at a constant annual rate until the end of the horizon. Data to parameterize the distribution for this uncertainty come from a representative sample of 14 Western utility IRPs, which make assumptions about the trajectory of carbon taxes over multidecadal planning periods.[31] All 14 utilities in this sample assume in their planning that the climate policy will come in the form of a tax.[32] If the model here is to reflect the beliefs of the decision-makers, it must use the metric that utilities incorporate into their own analyses.

Figure 3.2 plots the assumed CO_2 price trajectories over the planning horizons of the 14 Western utilities considered in this analysis. The reference trajectories are shown as solid lines, and the dashed lines represent other scenarios that are considered by utilities. These values are standardized into units of 2010$ per metric ton of CO_2e to facilitate comparison and integration into the capacity planning model.[33] The highly variable durations over which these trajectories span is indicative both of the range of publication years for utility IRPs and of the differences in planning horizons.[34] CO_2 prices and their escalation rates over time vary among utilities.

[31] The utilities selected for this analysis include: Avista, Idaho Power, Los Angeles Department of Water and Power (LADWP), NorthWestern, NV Energy, Pacific Gas and Electric (PG&E), PacifiCorp, Portland General Electric (PGE), Public Service Company of Colorado (PSCo), Puget Sound Electric (PSE), San Diego Gas and Electric (SDG&E), Seattle City Light, Southern California Edison (SCE), and Tri-State Generation and Transmission.

[32] Regulatory requirements for utilities like Idaho Power dictate that their IRP analysis of potential climate policies be performed as a carbon adder or tax.

[33] Short tons were assumed when IRPs did not explicitly specify the units.

[34] Planning horizons range from 10 years (PG&E, PacifiCorp, SDG&E, SCE) to 40 years (PSCo) with a mean of about 19.

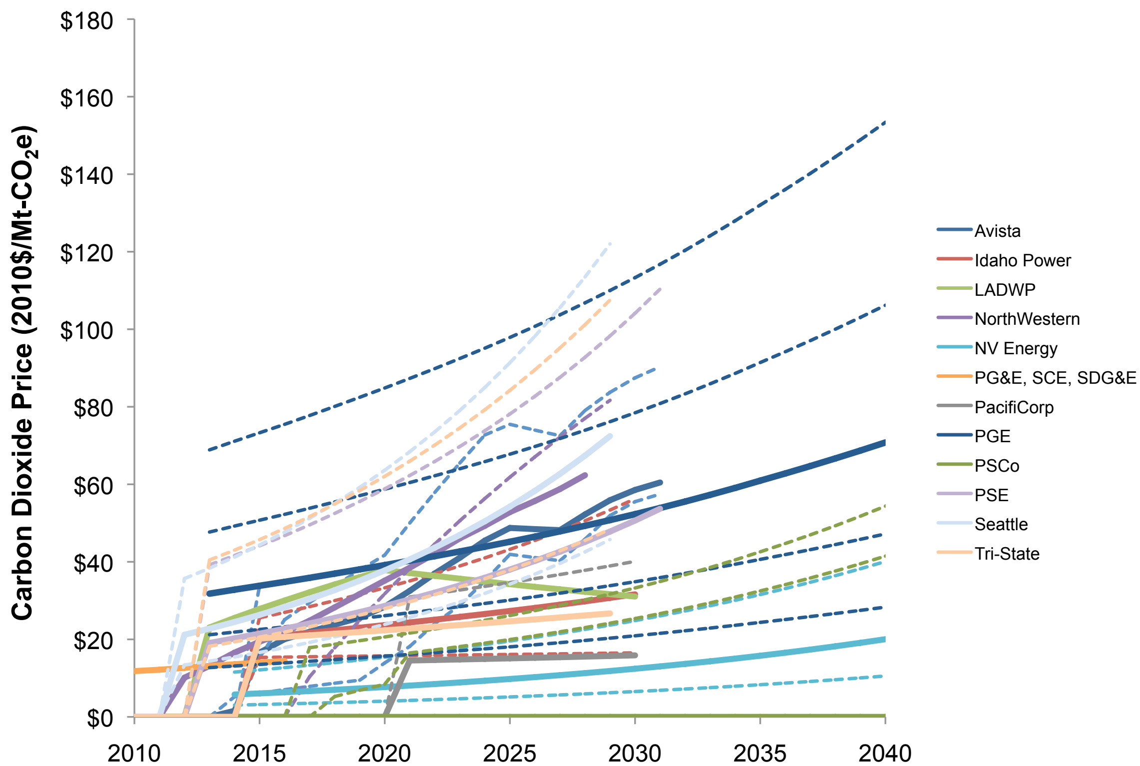

Figure 3.2: CO_2 price trajectory assumptions for 14 Western utilities in units of 2010$ per metric ton (Mt-CO_2e). The solid lines represent expected cases, and the dashed lines show alternate policy scenarios.

To illustrate the ranges of values considered in individual IRPs, Figure 3.3 presents the above data for the year 2025 only, which is the base model period when uncertainty is assumed to resolve. The reference or best-estimate cases are shown along with the ranges for the lowest and highest values considered by utilities. First, note that the reference case does not always correspond to the expected value of the carbon tax distribution, no matter how probabilities are assigned.[35] Second, many utilities view a no-policy scenario to be a serious possibility. 8 of the 14 utilities (57 percent) consider a no-policy scenario throughout the time horizon, and one utility (PSCo) uses the zero-valued tax as its reference value. Finally, comparing these ranges to estimates of the 2025 stringency of the Waxman-Markey bill illustrates that utilities'

[35]Only two utilities (Avista and NorthWestern) attach probabilities to the price scenarios.

expect that an implemented climate policy will have a much lower price on carbon. According to the Energy Information Administration (EIA) analysis of H.R. 2454 (DOE/EIA, 2009), the reference case ("ACESA Basic Case") indicates a 2025 price of about \$47/Mt-CO_2e. Only 5 of 14 utilities (36 percent) include this value in their potential range.[36]

The degree to which uncertainty analysis is incorporated into integrated resource planning varies across utilities. Only 2 of the 14 utilities (14 percent) attach probabilities to the uncertain climate policy scenarios. Ignoring probabilistic information precludes many of the uncertainty analysis and risk management approaches described in Chapter 2, including the ability to develop hedging strategies through sequential decision-making. This limitation is reflected in the common use of sensitivity analysis (79 percent), integrated risk metrics (43 percent), and threshold analysis (36 percent); however, no utility in the sample uses a sequential decision-making framework or uncertainty evaluation metrics.[37] Additionally, transparency about the treatment of risk is lacking in the resource plans, as only one utility explicitly defines how tradeoffs are made between expected cost and risk.[38]

[36] This finding is consistent with empirical evidence from cap-and-trade programs for sulfur dioxide and nitrogen oxide where, prior to the start of the programs, analysts overestimated the compliance costs of the proposed regulations (Taylor, 2012).

[37] Wilkerson, Larsen, and Barbose (2013) provide a detailed discussion of the types of risks considered by Western utilities in their long-term resource plans. They find that the risks most commonly included are future demand, fuel prices, and regulatory compliance for future climate policies.

[38] Barbose et al. (2008) reach similar conclusions about the treatment of carbon-price risk in utility resource planning and conclude that it is unclear if (or how) uncertainty about climate policy influences the selection of utilities' preferred portfolios.

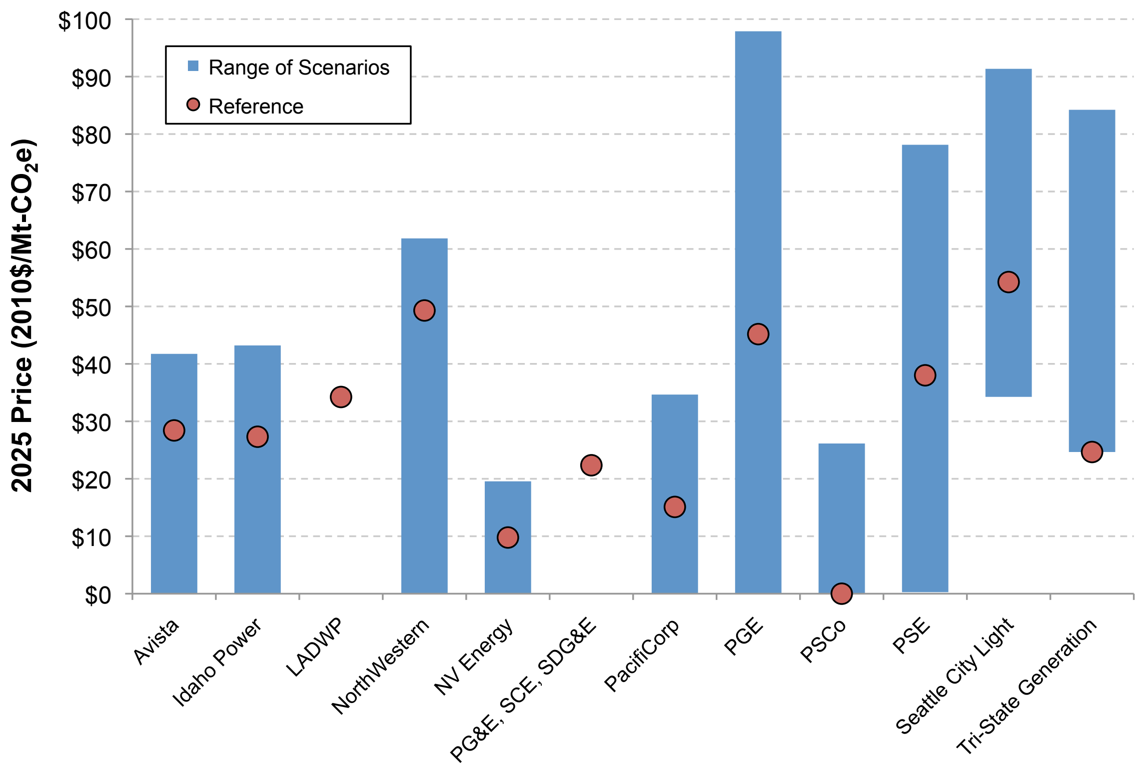

Figure 3.3: Carbon dioxide tax assumptions (2010\$/Mt-CO_2e) in 2025 by utility.

The CO_2 price trajectory data are used to create a distribution over potential 2025 prices in the model. Each scenario from each utility is treated as an independent draw from this distribution, and a five-point probability mass function is used to represent this discrete random variable, which preserves the sample mean and variance. The resulting probability density and cumulative distribution functions are shown in Figure 3.4, which has a mean of approximately \$30/Mt-CO_2e.[39]

[39]This research is agnostic about the methods used to arrive at the scenarios and distributions for carbon tax assumptions by utilities. These values must be taken as given, since there is no normative model for deriving such distributions, which differs from other uncertainties in the analysis.

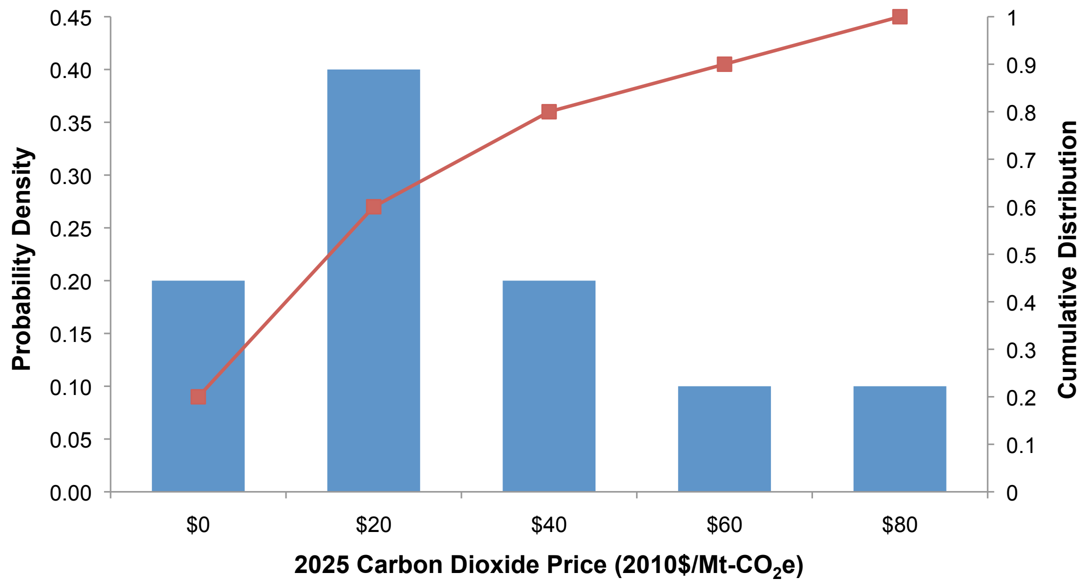

Figure 3.4: Discretized probability density and cumulative distribution functions for the second-stage CO$_2$e price (2010$/Mt-CO$_2$e).

3.4.2 Natural Gas and Coal Prices

Prices for energy resources are uncertain and fluctuate based on many complex factors. Uncertainty about the future of natural gas is also driven by recent discoveries and increased domestic production of shale gas (Moniz, Jacoby, and Meggs, 2010). Although abundant gas resources suggest expanded use in the electricity sector, uncertainty about long-run production costs and the environmental impacts of production make the extent of this growth unclear (Huntington, 2013; IEA, 2012a; DOE/EIA, 2011a; Coleman et al., 2011). Additionally, natural gas price uncertainty will be influenced by the unknown policy environment, public acceptance of hydraulic fracturing (Kriesky et al., 2013; Brown et al., 2013; Brasier et al., 2011), and uncertainty surrounding life-cycle emissions for shale gas (Howarth, Santoro, and Ingraffea, 2011; Jiang et al., 2011).

A vector autoregressive (VAR) model was created to estimate the probability distributions for future natural gas and coal prices of electric power generators. Using historical data for delivered fuel prices from the 2011 Annual Energy Review

(DOE/EIA, 2011b) published by the EIA and forecast data from the 2012 Annual Energy Outlook (DOE/EIA, 2012), this work uses a two-step process to estimate the trend and variability for future fuel prices and then uses this VAR model to create density functions for annual price growth rates for natural gas and coal. This model is based on the techniques developed by Zdybel and Baker (2013). Moment matching was used to discretize the resulting distribution into three-point discrete probability distributions for natural gas and coal. Appendix A discusses the VAR model specification in greater detail.

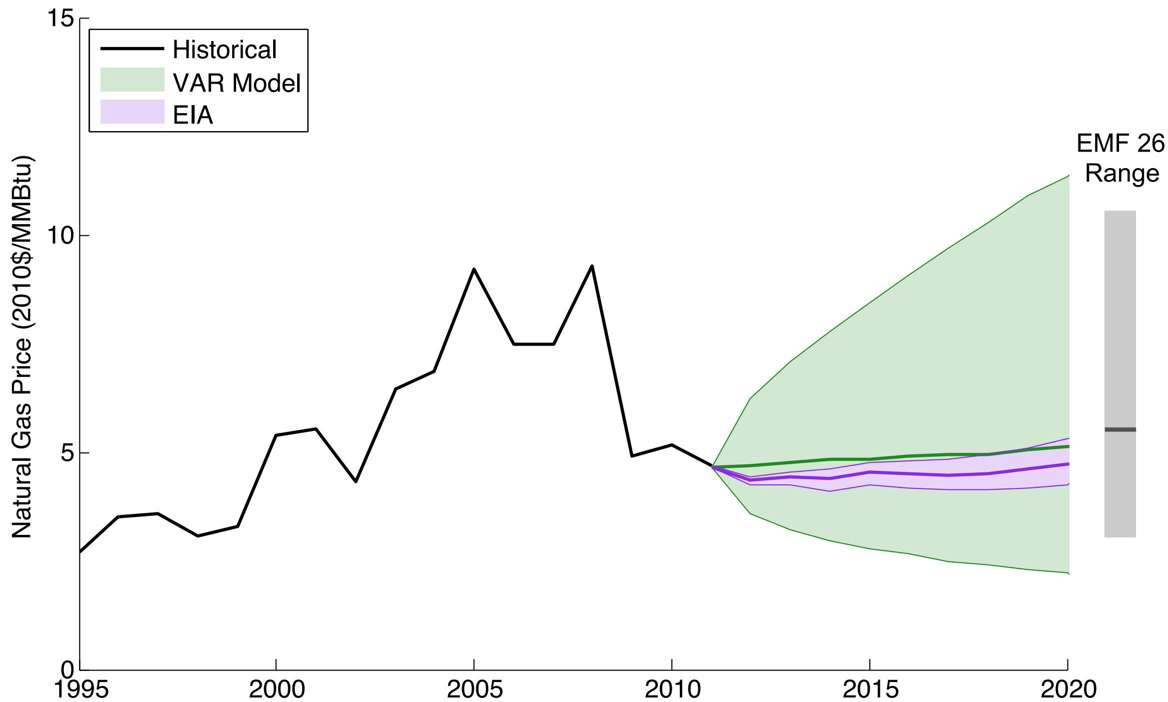

Figure 3.5: Historical and forecast delivered prices of natural gas for the electric power sector. The VAR model results show the 10^{th}, 50^{th}, and 90^{th} percentile values. The EIA cases represent the high-shale, reference, and low-shale scenarios. The EMF 26 cases represent the highest and lowest reported values in 2020 along with the reference scenario (averaged over all models).

Figure 3.5 shows the uncertainty ranges for delivered natural gas prices for the VAR model results and the 2012 Annual Energy Outlook (DOE/EIA, 2012). The trend for gas prices closely mirrors the EIA forecast, as the VAR model suggests that prices will increase only slightly over the next couple decades. However, the model results suggest that the uncertain range of prices may be much wider than the EIA projections, both on the lower and higher ends of the distribution. The VAR projections by 2020 are more consistent with the computational experiments from Stanford's Energy Modeling Forum (EMF) 26 study (Huntington, 2013), as shown in the bar at the right-hand side of the figure.

Modelers do not often quantify distributions over critical outputs or attach probabilities to possible scenarios, and there is evidence that, when analysts do quantify uncertainty, they tend to underestimate the range and probabilities associated with non-expected-value outcomes (Shlyakhtera et al., 1994). Although the EIA scenarios do not have associated probabilities, these results seem to support this finding and suggest that the overconfidence effect (i.e., the cognitive bias where confidence intervals are assessed too narrowly) also exists at an institutional level.

Based on this analysis, the annual natural gas price growth rate can take values of -5, 0, or 7 percent with probabilities of 0.30, 0.34, and 0.36, respectively. Coal price growth rates can have possible realizations of -2, 0, or 2 percent with corresponding probabilities of 0.28, 0.51, and 0.21, respectively.

3.4.3 Methane Emissions from Shale Gas Production

In addition to future prices, one of the most contentious and uncertain issues involving unconventional natural gas centers on the greenhouse gas impacts from its development. Research on life-cycle emissions from shale gas production has only been undertaken in the past two years, which focuses primarily on upstream CH_4 emissions. These studies exhibit a high degree of variation due to divergent assumptions and considerable uncertainty in the underlying data (Burnham et al., 2012; Cathles, Brown, and Taam, 2012; Howarth, Santoro, and Ingraffea, 2012; EPA, 2011a; Howarth, Santoro, and Ingraffea, 2011; Hultman et al., 2011; Jiang et al., 2011; DOE/NETL, 2011).[40] These problems are compounded by empirical data scarcities

[40] Across the life-cycle stages for natural gas systems, estimates of CH_4 emissions exhibit the greatest variation in production-related emissions due to assumptions regarding the frequency of

and the heterogeneity of sites and drilling practices. The lack of recent, direct emissions measurements combined with the tens of thousands of geographically diverse wells across the US will likely make this debate an active one in future years.

Figure 3.6 illustrates the disagreement and uncertainty in estimates of fugitive CH_4 emissions from shale gas across existing studies. This uncertainty is incorporated in the model as a random variable for emissions from shale gas production.[41] This work uses these values in a discrete three-point distribution with outcomes of 0.11, 0.6, and 1.18 grams of carbon per megajoule of fuel, which are interpreted as the 10^{th}, 50^{th}, and 90^{th} percentiles. Using the Extended Swanson-Megill approximation and assigning probabilities of 0.3, 0.4, and 0.3, respectively, to these outcomes provides a fairly robust approximation for a wide range of distributions (Keefer, 1994).[42]

On one hand, the uncertainty of prospective CH_4 leakage is more endogenous than other uncertainties, since producers can implement control technologies to reduce leakage rates. On the other hand, the upstream emissions rate is more uncertain than other random parameters considered in the analysis, since it is extraordinarily difficult to resolve unknown levels of past CH_4 emissions due to a lack of reliable measurements. Previous CH_4 emissions can be considered sunk from the firm's standpoint and do not impact future decisions. However, this uncertainty is much more problematic from a societal perspective due both to direct risks associated with climate damages and to indications about future monitoring credibility.

hydraulic fracturing to stimulate wells and of liquids unloading (Bradbury et al., 2013). The exceptions are the high emissions estimates from flowback and transmission in Howarth, Santoro, and Ingraffea (2012), which is one of the only estimates that does not rely primarily on EPA data.

[41] Upstream CO_2 and CH_4 emissions are included for coal and conventional natural gas production in the model as well.

[42] A recent study (Pétron et al., 2012) is one of the first to use actual air samples to characterize CH_4 emissions from shale gas systems. Using daily samples from the NOAA Boulder Atmospheric Observatory in Colorado, the multi-species analysis estimates that gas production in the Denver-Julesburg Basin leaks CH_4 at a rate that is twice as high as the Howarth, Santoro, and Ingraffea (2012) estimates for wellhead completion and production. The analysis does not quantify CH_4 emissions from other stages like leaks during distribution. Levi (2012) questions these findings, and Pétron et al. (2013) defend their initial results. A more recent evaluation (Karion et al., 2013) using aircraft overflight measurements for a gas field in Uintah County, Utah indicates an even higher CH_4 leakage rate, though the basin's emission rate may be atypical of surrounding regions. Results from an on-the-ground study (Allen et al., 2013) of onshore sites in the US are consistent with EPA estimates for production-related emissions, though potential selection effects may distort the representativeness of these results.

Although leaks represent only a few percent of the lifetime production of a well, CH_4 is the dominant portion of natural gas and a potent greenhouse gas, which means that even small leaks of this short-lived climate forcer can be significant. Recent modeling efforts (Shindell et al., 2009) have suggested that CH_4 may have an even larger global warming potential than previous estimates suggested (IPCC, 2007), particularly when indirect effects on atmospheric aerosols are taken into account. This analysis uses estimates of the global warming potential from Shindell et al. (2009) and a 100-year timescale to analyze the impact of CH_4.

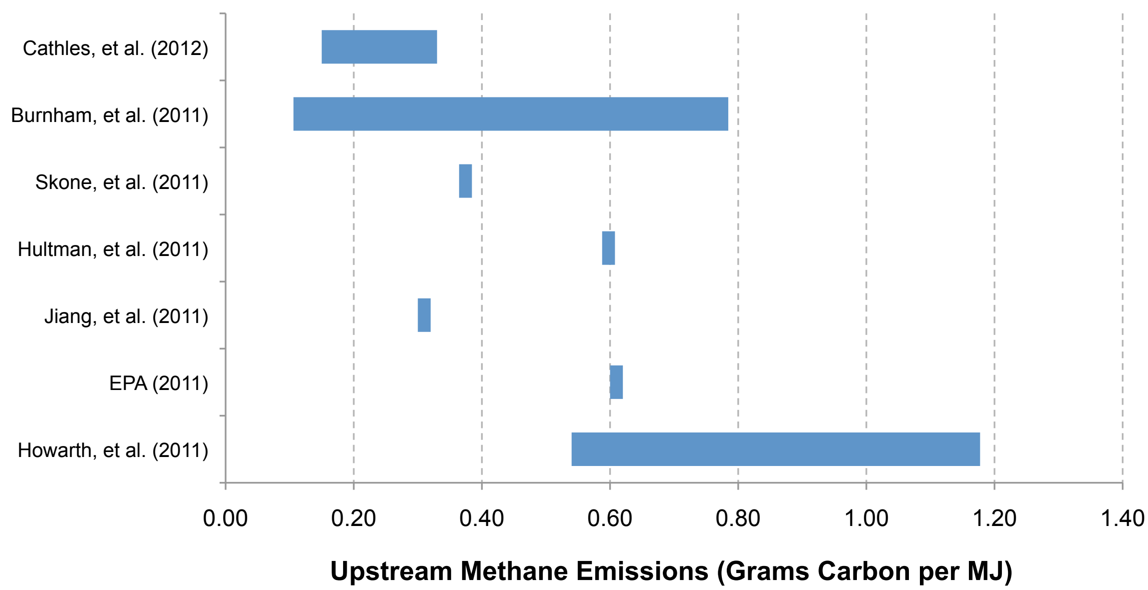

Figure 3.6: Estimates of fugitive methane emissions from shale gas production from the literature.

3.4.4 Capital Costs

There are many technological, economic, and geopolitical uncertainties that may impact CCS availability. These uncertainties include the effectiveness of CO_2 capture, possibilities of retrofitting plants and employing partial (or flexible) capture, unknown timeline for availability at scale, economics of building and operating CCS facilities, and presence of a politically viable policy framework for storing CO_2 (IPCC, 2005). This nascent technology is especially uncertain, as there are currently only eight CCS facilities worldwide and eight large-scale integrated projects were cancelled or postponed in 2012 (Global CCS Institute, 2012).

Similarly, nuclear power poses financial risks due to uncertainties in capital costs, regulatory requirements, electricity demand, and public opposition. Although these risks have been known for decades (OTA, 1984), the 2011 Fukushima Daiichi accident refocused attention on the perceived risks surrounding nuclear power. There is now uncertainty on how this post-Fukushima questioning of nuclear policies will impact the development and deployment of new reactors and retirement of existing capacity. Additionally, capital costs for new nuclear facilities historically have been higher than estimated, which complicates prospective forecasts. More than 75 percent of US nuclear reactors cost at least twice as much as initially estimated (DOE/EIA, 1986), and the French scale-up of reactors was characterized by negative learning, whereby construction costs increased with cumulative capacity installations (Grübler, 2010).[43]

Technological uncertainties for coal with carbon capture and nuclear are incorporated in the model as distributions over investment costs for these technologies. These distributions come from expert elicitations conducted at the Harvard Kennedy School (Anadon et al., 2011). Elicitation data are given as 10$^{\text{th}}$, 50$^{\text{th}}$, and 90$^{\text{th}}$ percentiles for experts in each technological area under base (business-as-usual) and enhanced R&D conditions.[44] The experts' elicited values are combined using Monte Carlo simulations weighting experts evenly. As with the original Harvard research (Anadon et al., 2011), the percentiles are fit to shifted log-logistic distributions. The resulting cumulative distribution functions are shown in Figure 3.7.

[43] Relative to other types of power plants, nuclear reactors are significantly more complex in their designs and operating conditions, which increases the potential for novel and unanticipated errors. For instance, a 1969 inspection revealed that abnormally high levels of radioactivity in the plant's drinking fountains were caused by a 3,000-gallon radioactive waste tank being improperly connected to the drinking water system. The Atomic Energy Commission's description understatedly reported, "The coupling of a contaminated system with a potable water system is considered poor practice in general" (AEC, 1969).

[44] The enhanced R&D distributions are used in the R&D strategy work in Chapter 6.

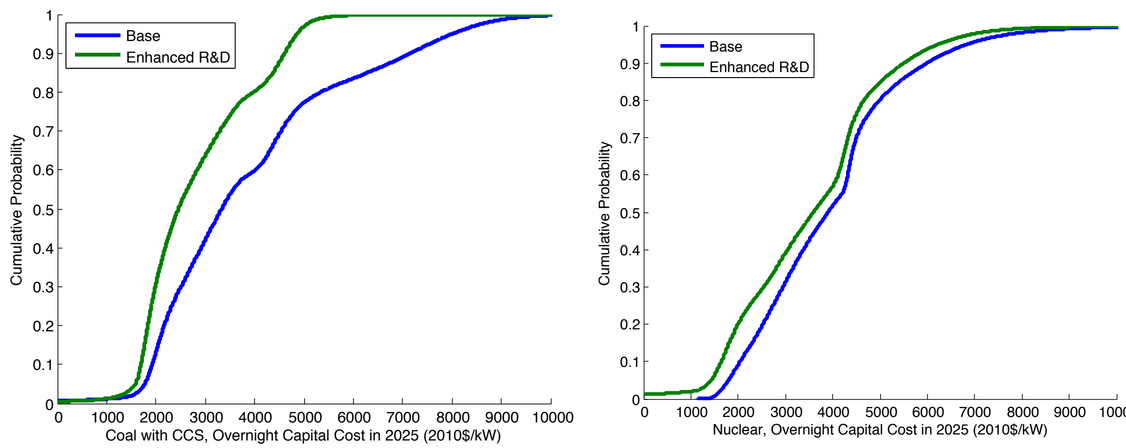

(a) Coal with CCS (b) Nuclear

Figure 3.7: Cumulative distribution functions for overnight capital costs.

For coal with CCS, using moment matching to create three-point discretizations yields potential outcomes of $2k, 4k, and 6k (2010$/kW) with probabilities 0.39, 0.44, and 0.17, respectively. For Generation III/III+ nuclear reactors, the random variable can take the value of $2k, 4k, and 7k (2010$/kW) with probabilities 0.28, 0.55, and 0.17, respectively.

3.4.5 Public Acceptance of CO_2 Storage

Although the technological, regulatory, and economic barriers to CCS development are considerable, public acceptance of CCS technologies and large-scale storage may be an equally daunting challenge.[45] Given questions about whether captured CO_2 will remain isolated from the atmosphere for long periods, it is not surprising that public concerns about CCS center on its environmental integrity (Einsiedel et al., 2013).[46]

[45] Despite greater public concern for climate change in recent years, only five percent of a 2007 surveyed sample had heard of CCS or carbon sequestration, and even these respondents had difficulties identifying which environmental problem CCS technologies address (Curry, Ansolabehere, and Herzog, 2007). This concern is especially pressing in the US relative to other countries like Sweden and Japan, where CCS awareness is closer to 15 and 22 percent, respectively (Reiner et al., 2006).

[46] Survey results in the US indicate that perception of CCS ranges from negative (Reiner et al., 2006; Palmgren et al., 2004) to slightly positive (Fleishman, Bruine de Bruin, and Morgan, 2010;

Public opposition has already been a factor in cancellations of proposed CO_2 storage projects (Van Noorden, 2010).

Putative risks like those surrounding large-scale CO_2 storage can enter into the decision-making process either directly (i.e., where outcomes have direct financial impact on decision-makers) or indirectly (i.e., where risk perception guides actions of involved stakeholders, and these actions have indirect financial impacts through channels like legislation or regulation). Thus, uncertainties like the public acceptance of CO_2 storage may be considered in the decision-maker's analysis based on perceived risk regardless of actual risk, unless these two can be bridged through the intervention of researchers, journalists, and decision-makers before such risks acquire culturally divisive meanings (Kahan et al., 2012).[47]

The public acceptance uncertainty for large-scale CO_2 storage uses probability estimates from a National Academies study (National Research Council, 2007), which accounted for public opposition based on risks from sequestration and siting requirements. According to this study, the probability of normal storage is 0.66, which implies a 0.34 probability that storage is prohibited.[48]

3.4.6 Natural Gas Combined Cycle Efficiency

In addition to uncertainties about the economic and environmental impacts of unconventional natural gas, another relevant uncertainty that will shape the future role of natural gas in the electric power sector is the performance of gas-turbine-based

de Best-Waldhober, Daamen, and Faaij, 2009). Acceptance of CCS ranked lower than nuclear power in some pre-Fukushima surveys of public acceptance (Reiner et al., 2006). Support for CCS has been shown to increase when information on costs and environmental benefits of CCS are provided in some studies (de Best-Waldhober, Daamen, and Faaij, 2009; Reiner et al., 2006), but other surveys have indicated declining support when information is given (Palmgren et al., 2004).

[47]Research suggests that acceptance and support for CCS may be influenced by beliefs about climate changes (Einsiedel et al., 2013). This link may make it challenging for educational campaigns to avoid attaching antagonistic cultural meanings to CCS, since it is difficult to convey information about CCS and its benefits without mentioning the technology's link to climate change.

[48]Public opposition likely will not manifest itself in a binary, all-or-nothing outcome for utilities across the country. A more realistic modeling approach would have more possible manifestations of these restrictions, especially since regulations and risk assessments are likely to come at a state level (similar to hydraulic fracturing). The current framing also could be conceptualized as the long-term technological success of integrated carbon capture projects.

technologies. In particular, the first-law efficiencies of these technologies with and without carbon capture may determine the diffusion of new capacity and market share of generation from natural gas. Such characteristics are especially important for a technology subject to fuel price volatility and to similar levelized electricity costs as other technological substitutes, which means that even small efficiency changes may have large impacts on future diffusion and utilization.

I conduct a series of elicitations with experts in industry and academia to investigate the best practices of energy technology probability assessments through a case study of a policy-relevant technology that has been hitherto neglected in the elicitation literature. In particular, the aim of this work is to represent the current state of knowledge regarding the future of gas turbine systems for new central station electricity generation. Most elicitations for fossil-based electricity generation technologies focus on coal with CCS, and when research groups look at gas with CCS, it is typically to encode uncertainty about capital costs. Here, I elicit expert judgments on the first-law efficiencies of commercially viable natural-gas-fired power plants.

In the absence of this elicitation approach, most energy-economic models simply assume that future plant efficiencies will remain constant at current levels (with combined cycle efficiencies between 50 and 60 percent) or will marginally increase between now and 2050. Even slight deviations from these efficiency values can have large impacts over time in the development and deployment of gas-turbine-based systems, particularly when natural gas prices and climate policy are uncertain and there are many substitute technologies and fuels (as discussed in Chapter 4.2).

Figure 3.8 shows the cumulative distribution function of elicited values for first-law efficiencies in 2025 under the business-as-usual R&D scenario. Individual values for all four experts are given along with the combined and fitted cumulative distribution function. Although the figure shows some disagreement, particularly for higher efficiencies, it is notable that all experts agree that the median efficiency value for **2025** will be at least 60 percent. According to the analysis in Appendix B, only one existing energy-economic model has an efficiency value that exceeds 60 percent through **2050**.[49] Thus, existing models significantly underestimate performance characteristics for future natural gas systems for electricity generation.

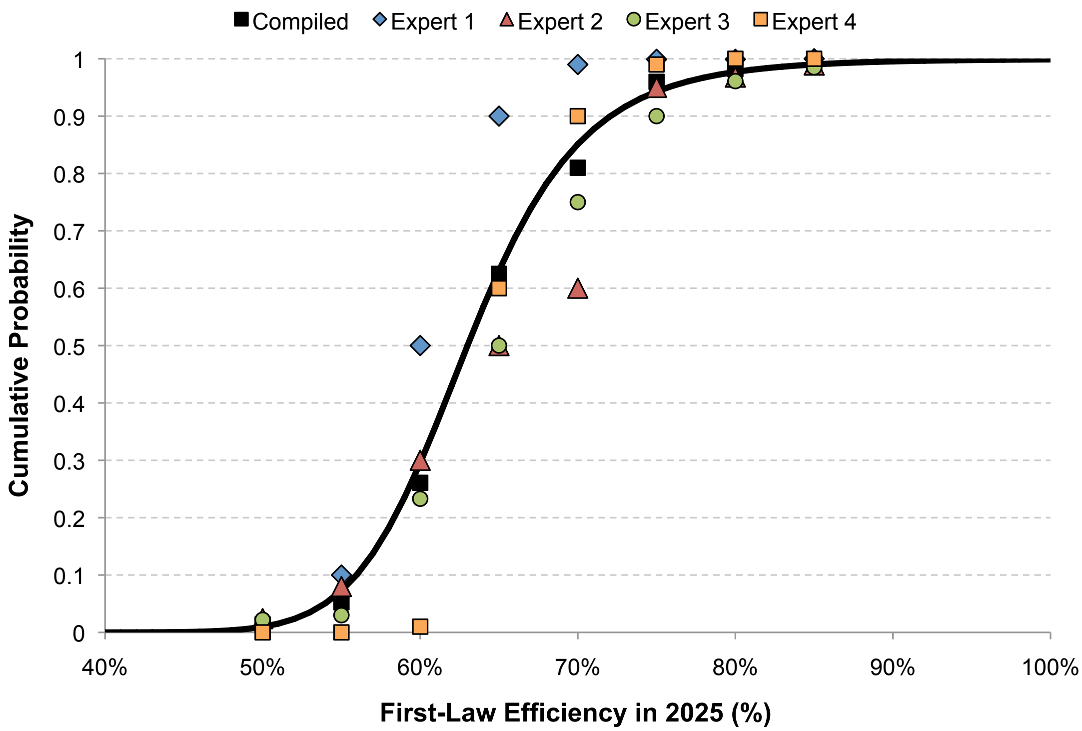

Figure 3.8: Elicited values for first-law efficiencies (lower heating value basis) of new gas-turbine-based electricity generators in 2025.

Using the Extended Swanson-Megill approximation, the three-point discretization of this random variable can take the value of 56, 63, and 72 percent with probabilities 0.3, 0.4, and 0.3, respectively.

A more in-depth discussion of the elicitation protocol and results is found in Appendix B.

[49]The Siemens SGT5-8000H gas turbine achieved a world record 60.75 percent efficiency in a combined-cycle configuration at the Irsching Power Station in Bavaria, Germany in May 2011.

3.4.7 Model Representation of Uncertainties

Table 3.2: Uncertainties included in the capacity planning model.

Uncertainty	Scenarios
CO_2 tax stringency	5
Natural gas price path	3
Coal price path	3
Upstream CH_4 emissions	3
Coal with CCS capital costs	3
Nuclear capital costs	3
Public acceptance of CO_2 storage	2
Natural gas combined cycle efficiency	3
Total	*7,290*

Each of these random parameters is assumed to be statistically independent (i.e., probabilistically irrelevant).[50] This assumption means that there are a total of 7,290 universe scenarios, as shown in Table 3.2.

The model uses a two-stage stochastic programming approach and is programmed in the General Algebraic Modeling Software (GAMS) environment using the DECIS system (Infanger, 1999). All uncertainties are assumed to resolve completely in 2025. From this model period forward, second-stage decisions can be made with complete and perfect knowledge of all future parameters.

[50]The specifications of random variables are made so that these uncertainties are orthogonal as much as possible, acknowledging that characterizing marginal distributions for uncertain parameters is difficult enough let alone the complete joint distribution. For instance, a capital cost inflation metric over time is used in Chapter 4.2.3 to model correlated construction costs for different plant types. Although the uncertainties have been defined and selected in a manner to reduce the potential for correlation, future research efforts should attempt to more rigorously quantify potential impacts of correlations between random variables. For example, natural gas prices and CH_4 leakage rates may be negatively correlated if low prices provide strong incentives to reduce costs by eliminating control technologies and other practices that could have reduced emissions during well production and completion.

Chapter 4

Capacity Planning Results

4.1 Reference Results

Table 4.1 lists objective function values for the wait-and-see (z_{ws}), stochastic (z^*), and expected-value (z_d) solutions.[1] The top rows list values when the uncertainties are considered individually, and the bottom row shows results for all eight uncertainties considered jointly.

The joint expected value of perfect information (EVPI) of $162 billion is considerably larger than the value of the stochastic solution (VSS) of $36 billion,[2] which indicates that resolving uncertainty is more valuable than simply accounting for it in modeling efforts and implementing a hedging strategy.[3] Additionally, the joint EVPI and VSS come primarily from three uncertainties: climate policy, natural gas prices, and public acceptance of large-scale carbon dioxide (CO_2) storage. The highest values correspond to the natural gas price uncertainty, which underscores the significance of this factor in utilities' capacity-planning decisions.

[1]The numerical results in this section should be interpreted in the context of the accompanying model assumptions from Chapter 3. Greater emphasis should be placed on the insights gleaned from this framework rather than on the exact magnitudes of the model outputs.

[2]All monetary values are expressed in US 2010 dollars with a discount rate of five percent unless otherwise noted.

[3]The result that the value of information acquisition exceeds the value of precautionary investments mirrors the conclusions of Manne and Richels (1993).

Table 4.1: Comparison of discounted system costs (billion $) for the wait-and-see ($z_{ws}$), stochastic ($z^*$), and expected-value ($z_d$) solutions when uncertainties are considered individually and jointly. The EVPI and VSS metrics are shown in the rightmost columns, respectively.

Uncertainty	z_{ws}	z^*	z_d	**EVPI**	**VSS**
Stringency of abatement policy	4,114	4,171	4,185	57	15
Natural gas prices	4,096	4,168	4,204	72	36
Coal prices	4,278	4,279	4,279	2	0
Upstream CH_4 emissions	4,283	4,283	4,283	0	0
Capital costs (coal with CCS)	4,283	4,283	4,283	0	0
Capital costs (nuclear)	4,283	4,283	4,283	0	0
Public acceptance of CO_2 storage	4,283	4,290	4,291	7	1
Natural gas combined cycle efficiency	4,283	4,283	4,283	0	0
Joint	3,884	4,047	4,083	162	36

4.1.1 Value of the Stochastic Solution

To account for the relatively low VSS, it is instructive to compare first-stage decision variables between the stochastic and expected-value solutions. Table 4.2 lists capacity investments by generator type before the uncertainties are resolved in 2025.[4] The similarity between the stochastic and expected-value strategies accounts for the small VSS. Capital investments are concentrated primarily in new wind and nuclear assets, which are common to both approaches. This result suggests that these technologies are strong candidates for near-term hedges against a variety of uncertainties while allowing the power sector to keep pace with growing demand and retirements of significant portions of the current generator fleet in the coming decades. These technologies are attractive investments due to their low life-cycle greenhouse gas emissions (which reduces exposure to the climate policy uncertainty) and to their low fuel price volatility (relative to alternatives like natural gas), which means that they are economical under a wide range of contingencies and are not as likely to be mothballed or decommissioned once new information becomes available.

[4] Values in the table represent planned capacity but not necessarily completed additions by 2025. Construction lead times in the model mean that nuclear units will not come online until two periods (i.e., ten years) after the investment decision is made (and one period for CCS-equipped capacity).

Table 4.2: Cumulative capacity investments (GW) by 2025 under the stochastic and expected-value solutions.

	Stochastic	**Expected Value**
Biomass	3	3
Coal with CCS	0	61
Natural gas combined cycle	0	0
Nuclear	289	287
Wind	139	139
Total	*431*	*490*

The most striking difference between these strategies is the 61 GW investment in new capture-equipped coal generators under the expected-value strategy.[5] In part, the stochastic strategy avoids near-term carbon capture and storage (CCS) investments due to the possibility that these assets would be mothballed or decommissioned either if the climate policy is too stringent or too lax or if public opposition prevents cost-effective CO_2 storage. The stochastic strategy delays investment in new capacity and instead relies on increased generation from existing, underutilized (i.e., low capacity factor) natural gas units.

Apart from this stranded-cost effect, CCS investments under the stochastic approach are lower as a means to avoid irrevocably committing resources to capital assets that not only may be suboptimal for the realized state of the world but also may displace investments in more profitable generating capacity in later periods. Second-stage capital expenditures, made after more information about uncertain quantities is available to utilities, may profitably adjust deployment levels of technologies to take advantage of unforeseen and unlikely boons like unexpectedly low natural gas prices. In this instance, the stochastic strategy avoids irreversible investments early on, which may entail sizable opportunity costs from foreclosed opportunities, by not investing in CCS-equipped coal units and by instead meeting growing demand with increased utilization of existing natural gas units.[6] In contrast, the mean value of

[5]In addition to overestimating the amount of coal with CCS relative to the stochastic strategy, the expected-value solution also slightly underestimates the optimal amount of nuclear, since the stochastic solution is providing a greater hedge in the event that a moderate climate policy is realized but large-scale CCS is not available.

[6]Investments are considered irreversible when expenditures are sunk costs and it is impossible

the carbon tax distribution of \$30/Mt-CO_2 lies in the permissible policy range where coal with CCS is lucrative, which means that the expected-value approach deploys this technology during the first stage.

This insight about the small but important bifurcation between the stochastic and expected-value approaches is particularly interesting given the typical behavior of similar decision problems under uncertainty. Hedging policies, like the stochastic approach here, often diversify investments across a range of technological options to protect against failure (Birge and Rosa, 1996). In this instance, utilities choose to delay investments with the hope of learning more information instead of making additional precautionary investments. The option to postpone decisions gives utilities the ability to tailor capacity investments to the realized scenario, which makes this freedom to wait valuable if the opportunity costs of delay are relatively low (Dixit and Pindyck, 1994). The threat of stranded assets from irreversible investments results in delaying capital outlays. In general, the VSS is higher for uncertainties that induce anticipatory, near-term actions.[7]

The profitability of delaying a fraction of investment under uncertainty is an effective strategy for utilities in the United States (US) context for three reasons. First, as mentioned already, the stochastic strategy can increase generation from underutilized natural gas combined cycle (NGCC) capacity in the interim before uncertainty is revealed. Second, exogenous retirements due to plants exceeding their operating lifetimes will not occur *en masse* for a couple of decades. By that time, more information will likely be available about long-term policy trajectories, fuel prices, and technological cost and performance characteristics. Finally, slower projections of electricity demand growth in the coming years will obviate the need for new capital investments right away.

A considerable amount of variance for the cost advantage of the stochastic strategy under different states of the world can be accounted for by two of the most important

---

to uninvest, which occurs when economic conditions decline and investments are industry-specific (Tuthill, 2008).

[7]Delays also could result from risk aversion, which is not incorporated in this analysis; however, as noted by Arrow and Fischer (1974), "Something of the 'feel' of risk aversion is produced by a restriction on reversibility."

uncertainties, as illustrated by Figure 4.1. The decision not to build CCS capacity before uncertainties are resolved proves to be most valuable under scenarios with low natural gas prices and lenient climate policies. In these scenarios, the first-best option under uncertainty is to build large quantities of NGCC units during the second stage once information has been revealed about modest carbon and low gas prices. However, for the expected-value approach, irreversible investments in coal with CCS during the first stage would provide less flexibility for taking advantage of these market opportunities.

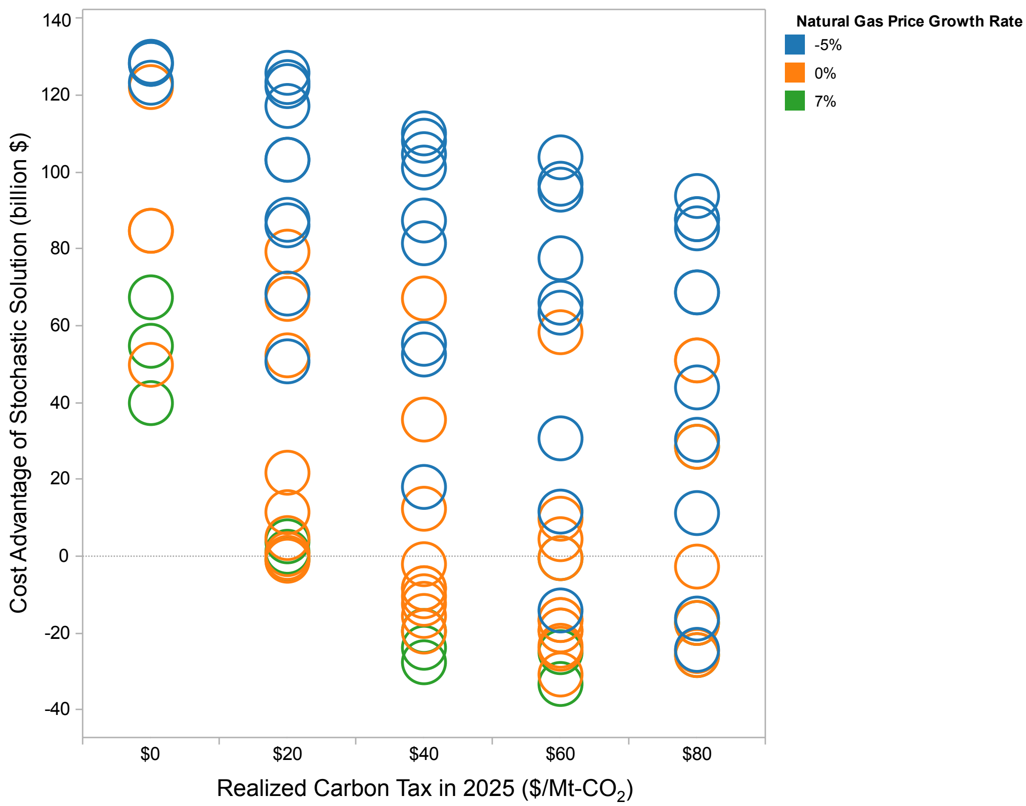

Figure 4.1: Cost advantage of the stochastic solution with different realized values of the climate policy uncertainty (horizontal axis) and natural gas price uncertainty (colors). Each circle represents one scenario of the 7,290 possible outcomes.

One reason why explicitly accounting for uncertainty leads to prescriptions of decreased near-term investments in low-carbon technologies relates to abatement cost characteristics of the power sector. Total system costs (i.e., the objective function in the utilities' intertemporal optimization problem) are concave in carbon taxes, with costs essentially plateauing around \$20/Mt-CO_2. This general trait is consistent with a wide range of studies in the literature that find that a carbon tax of \$14–27/Mt-CO_2 would largely decarbonize new electricity generation (Weyant, de la Chesnaye, and Blanford, 2006), which means that more stringent carbon taxes would not substantially impact costs toward the second half of the modeling horizon.

Consequently, if utilities plan for the mean carbon tax (\$30/Mt-CO_2) in the first stage and the most stringent case is later realized, higher carbon taxes would not substantially impact second-stage costs. Thus, the only method of achieving cost savings for the climate policy uncertainty is for utilities to avoid irreversible investments early on and hope that the low-tax scenario is realized so that second-stage investments are less costly. If one were to graph a loss function as in Figure 4.2 (with expected system costs on the vertical axis given the assumed carbon tax on the horizontal axis), the function would be minimized for a carbon tax between \$15–20/Mt-CO_2, which corresponds to the optimal stochastic strategy. The function increases steeply around a zero carbon tax, indicating that, although it is more beneficial to err on the lower side of the mean, it is costly to undershoot the optimal value by too much.[8]

[8]The importance of the actual distribution to this insight illustrates how empirically-derived distributions are required to resolve ambiguity about optimal energy decisions in uncertain environments (Baker and Solak, 2013). It shows the value of incorporating actual data instead of stylized, *ad-hoc* distributions.

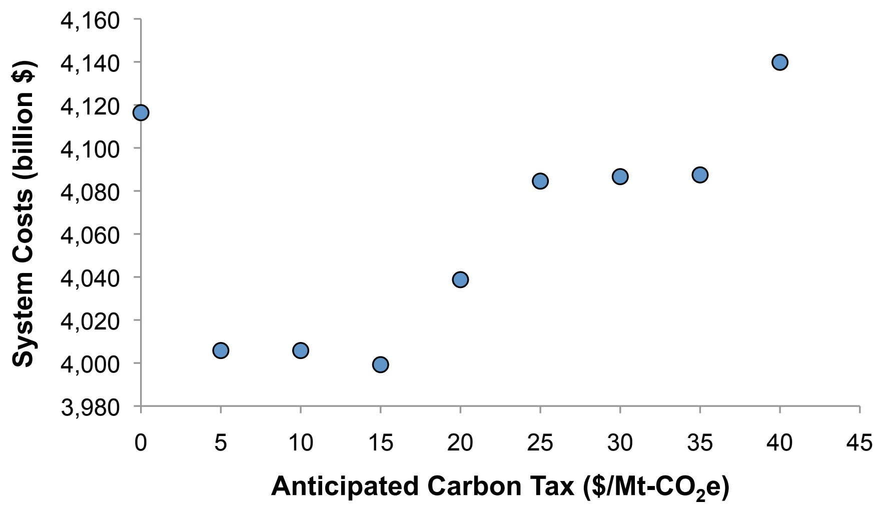

Figure 4.2: Loss function of expected system costs (vertical axis) for different first-stage planning assumptions about future CO_2e taxes (horizontal axis).

The conclusion that it is suboptimal for the stochastic solution to construct coal with CCS capacity during the first stage is a robust one. I tested the robustness of this recommendation to the number and types of uncertainties included in the analysis and found that, even when only the climate policy and gas price uncertainties are considered, the result still holds. When only the climate policy uncertainty is considered, there is a small amount of CCS-equipped coal generation built (12 GW), but it is a much smaller investment than the expected-value solution (61 GW).

In another experiment, I assume that utilities and generators plan based on a "best-guess" assumption of no substantial climate policy instead of using the mean value for the second-stage carbon tax. This no-policy scenario can be interpreted in several ways: 1. Expectation of stalled international negotiations and/or disagreement at the federal level; 2. Anticipation of a discovery that climate change is not as threatening as expected; 3. Expectation of a low-risk geoengineering solution or cheap ambient air capture to decouple emissions from climate impacts. As mentioned in Chapter 3, utilities may view climate policy as an endogenous uncertainty with (partially) controllable outcomes, particularly if climate risk is noisy and high evidentiary standards for control can delay action. Even if a carbon tax materializes, this scenario could represent the case where utilities and generators believe that there is a high likelihood that carbon-intensive units built before policy is enacted will be grandfathered into legislation (i.e., creating a self-fulfilling prophecy danger).[9]

If the VSS is computed assuming a no-policy baseline, the VSS increases to $61 billion, which is much higher than the $36 billion using the expected-value solution. This VSS can be interpreted as the expected cost of inaction. In contrast to the baseline case, the no-policy solution builds no nuclear capacity (instead of 287 GW) and instead constructs 107 GW of coal-fired units. These units would represent large financial losses if emissions restrictions are later put in place or if natural gas prices are exceedingly low, which would lead to a large-scale decommissioning of these units. These scenarios would cause these units to be decommissioned almost immediately, which gives rise to a larger VSS.

This experiment illustrates the importance of model assumptions about a decision-maker's expectations, why cancellations of planned coal additions observed around 2007 initially occurred (i.e., due to expectations that a carbon price would materialize in the near future), and why few coal additions have been proposed since (i.e., due to expectations that natural gas prices will remain low). This expected-value solution only performs well under a distribution that is compatible with the assumed no-policy prior, but the strategy is vulnerable when it encounters a world with a dramatically different distribution.

4.1.2 Expected Value of Perfect Information

Much like the VSS, the EVPI of $162 billion is driven primarily by information for scenarios with less stringent carbon taxes and lower natural gas prices, as shown in Figure 4.3. Information has value in lax climate policy scenarios, since it would be optimal to build fewer low-carbon units like wind and nuclear during the first stage and instead wait to build fossil-based capacity, especially if natural gas prices are low.

[9]The utilities' climate policy distribution ultimately hinges on epistemological and psychological questions related to belief formation and mental models of risk, which are beyond the scope of this research. Additionally, the utilities' optimization problem may be more complex than suggested by this formulation, as values may be excluded from the objective function or some parameters may be considered decision variables instead of given values.

The utility of having access to early information comes through the ability to avoid irreversible, capital-intensive investments in technologies that are suboptimal for the realized scenario.

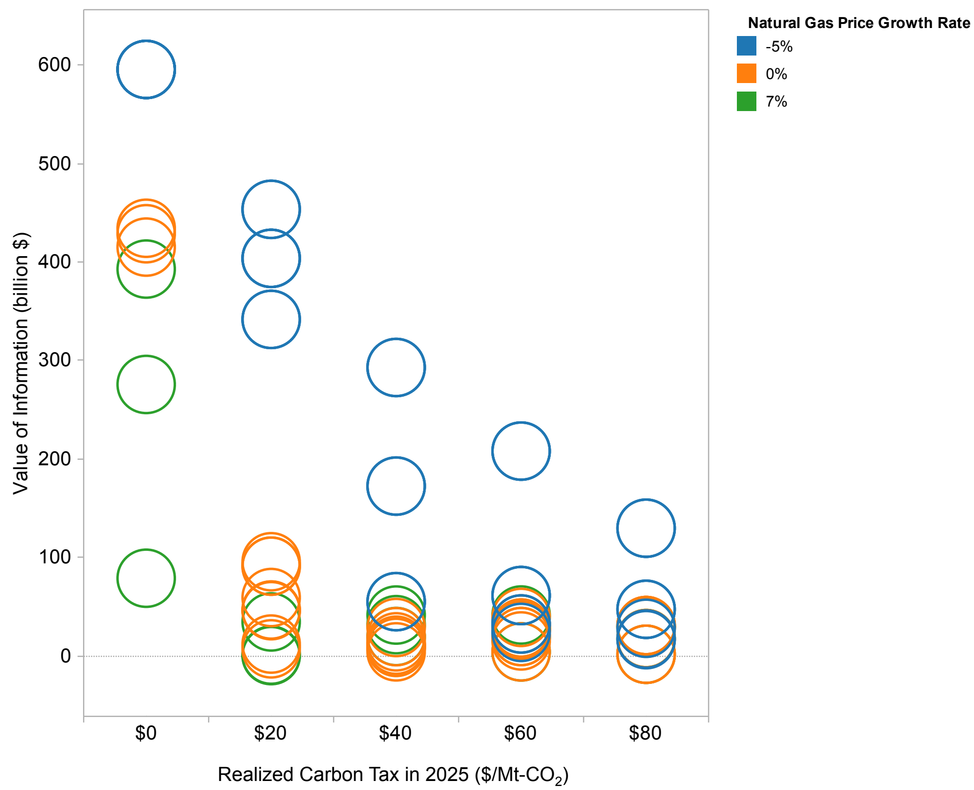

Figure 4.3: Value of information with different realized values of the climate policy uncertainty (horizontal axis) and natural gas price uncertainty (colors).

4.1.3 Shale Gas

Many of the most pressing energy policy questions in the US are related to unconventional natural gas resources like shale gas. The model runs in this section contain many insights about the potential role of shale gas as part of the domestic energy mix. Figure 4.4 suggests that shale gas resources are used for electricity generation largely when natural gas prices are low (with abatement targets being a secondary driver). Shale gas and natural gas in general are less important for ambitious climate targets no matter what price is assumed and regardless of upstream emissions. The dark green area in Figure 4.4 illustrates that only small amounts of shale gas are used when the natural gas price growth rate is at its mean value or higher.

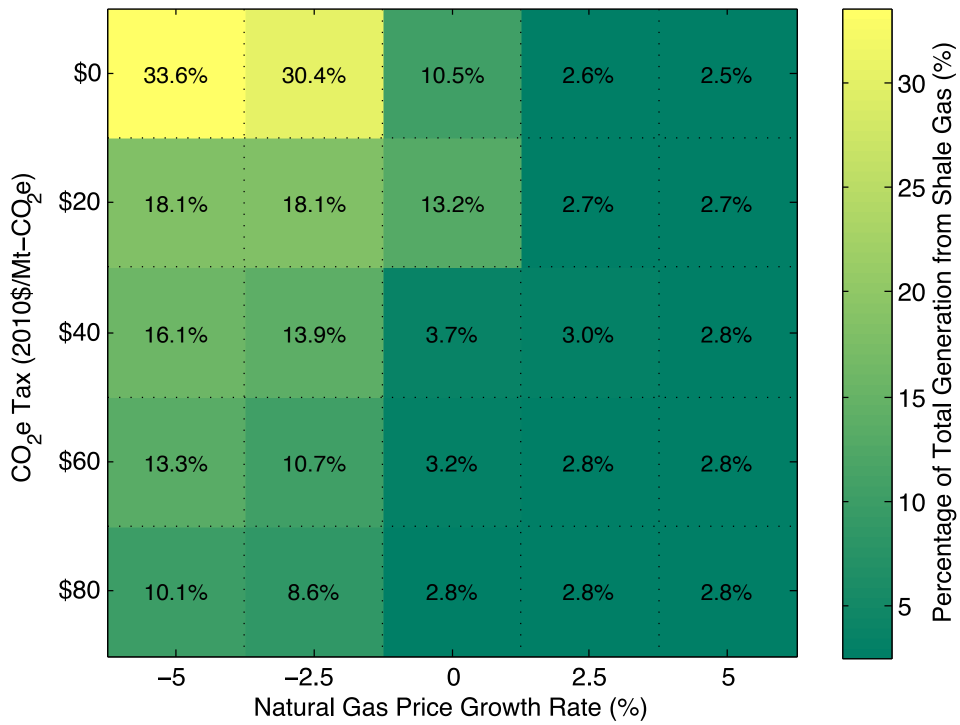

Figure 4.4: Percentage of total generation from shale gas after 2010.

Reference results for capacity investments in Table 4.2 indicate that there is no deployment of new natural gas capacity under the stochastic or expected-value strategies before uncertainty is reduced. In the next decade, lower natural gas prices are likely to spark greater utilization of existing capacity rather than new construction. The greatest potential for new capital investments will occur later in the future once more information is available and once a stronger long-term price signal can lower investment uncertainty.

Another related question is whether a carbon price will increase or decrease natural gas consumption in the power sector (Huntington, 2013). Figure 4.4 illustrates how the answer to this question varies depending on the realized gas price. When natural gas prices are low, the no-policy case relies primarily on gas generation, which means that natural gas consumption is monotonically decreasing in the carbon tax. However, when natural gas prices are close to the mean value of the distribution, the no-policy case involves generation with substantial amounts of both coal and natural gas. Therefore, when average gas prices obtain, increasing the carbon tax leads to more generation from gas-fired capacity up to a point and then begins to decrease.[10] This effect illustrates the importance of modeling the interactions between multiple uncertainties simultaneously, particularly for complex policy questions.

There is also widespread interest in determining how shale gas availability will influence investments in renewable technologies. Speculation centers on questions about the degree to which a low-cost shale boom may curtail the deployment of low-carbon substitutes like wind and nuclear. The ternary plot in Figure 4.5 illustrates that many more considerations than simply natural gas prices will influence how gas could displace investments in other electric sector technologies.[11]

The expansion of generation from natural gas units is largest under scenarios where gas prices are low and the stringency of climate policy is low to moderate. Under a scenario where no climate policy is enacted, generation comes primarily from fossil-based units, with gas comprising nearly 85 percent of generation by 2050 when gas prices are low. The availability of low-cost shale gas lowers greenhouse gas emissions by replacing production from coal, even though no climate policy is in place in this state of the world. When a moderate policy is enacted and abundant reserves lower gas prices, coal is eliminated from the generation mix by 2050, and 50 percent of electricity comes from natural gas.[12] For this specific case, the existence of low-cost shale gas means that gas units replace what would have otherwise been predominately coal with CCS (which would generate 21 percent of generation by 2050), nuclear, and

[10] For mean-valued gas prices, note that increasing the carbon tax means more natural gas generation on a relative basis (i.e., compared with the no-policy case). On an absolute basis, generation is still not as large as when natural gas is cheapest.

[11] Ternary plots use barycentric coordinate systems to depict proportions of three variables as locations on an equilateral triangle. The proportions of these three components sum to a constant value (typically 100 percent, as in Figure 4.5).

[12] Since natural gas is still a hydrocarbon-based fuel, approximately a third of the natural gas generation comes from CCS-equipped units in order to comply with the moderate climate policy.

wind. Under a stringent climate policy scenario, however, the presence of shale gas in the resource supply curve for natural gas has very little influence on the deployment of technologies in the power sector. The model generates a significant fraction of electricity from non-emitting resources like renewables and nuclear by 2025 regardless of the gas price. Thus, the influence of shale gas on electric sector investments depends strongly on the stringency of the climate policy in addition to natural gas prices.

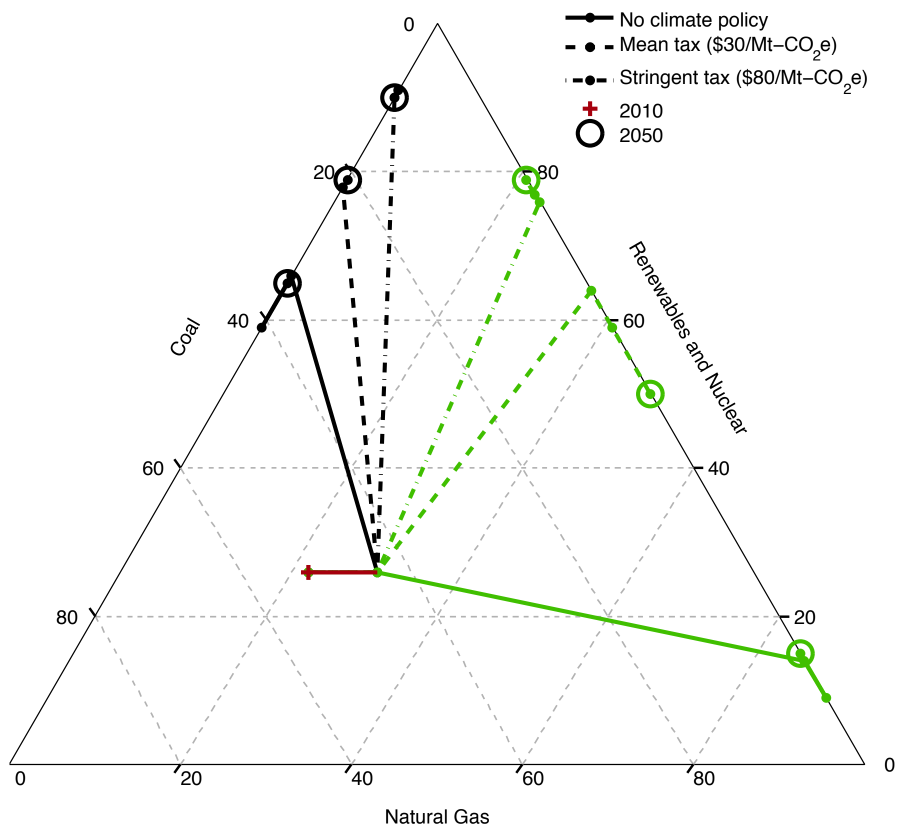

Figure 4.5: Ternary plot of electricity generation share (%) by technology under various gas price and climate policy scenarios, 2010–2050. The gridlines indicate the fraction of total generation in a given year from renewables and nuclear (horizontal gridlines), natural gas (diagonal gridlines from the lower left to upper right), and coal (diagonal gridlines from the upper left to lower right). High gas price scenarios are depicted in black and low price scenarios in green with 2010–2020 values shown in red. Note that, since the stochastic hedging approach is used, the strategies are the same for all scenarios before the uncertainty resolution date of 2025.

Another research question with significant policy dimensions is how much utilities and generators would be willing to pay for research, development, and deployment of control technologies to limit fugitive methane (CH_4) emissions from shale gas. This value of control places an upper bound on the deployment of control technologies and can be calculated by taking the difference between the expected cost of the stochastic strategy (with all 7,290 scenarios) and the expected cost of the problem where CH_4 leakage is certain to be zero.

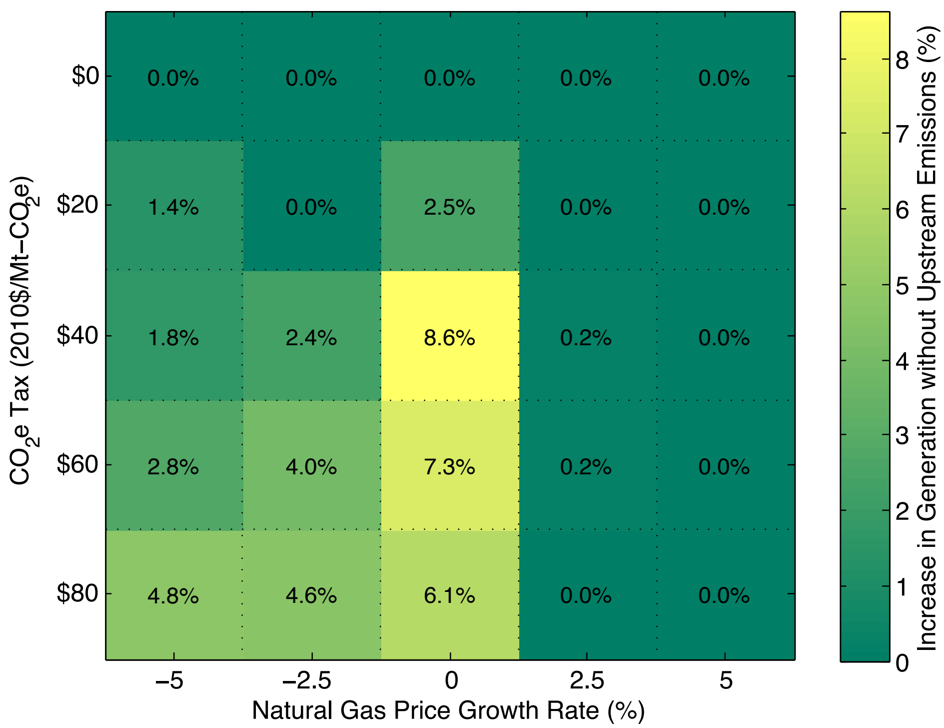

Figure 4.6: Percentage increase in generation from shale gas (on an absolute basis) when upstream methane emissions are zero instead of being uncertain.

The value of control is $40.5 billion, which indicates that there is considerable benefit to limiting upstream emissions. Limiting CH_4 leakage allows more gas units to be built and operate during the second stage in scenarios where higher carbon taxes are realized and natural gas prices are low to moderate, as shown in Figure 4.6.[13]

[13] The benefits of a CH_4 control technology would be even greater if a correlation exists between low gas prices and high leakage rates, which is not incorporated in the model.

The ability to limit upstream CH_4 emissions would be most beneficial in the parameter space where the carbon tax is about \$40/Mt-CO_2 and gas prices are relatively constant over time. Under such conditions, generation that would have come from coal with CCS (when upstream emissions are uncertain) instead would be replaced by natural gas with CCS. The magnitude of this substitution effect is shown in Figure 4.7 and results in an 8.6 percent increase in generation from shale gas, which takes advantage of the lower natural gas prices. Thus, limiting upstream emissions represents a large value-added proposition for utilities and shale gas developers, since it can allow greater deployment of gas units under climate policy and gas resource scenarios when they otherwise would not have been cost-competitive.

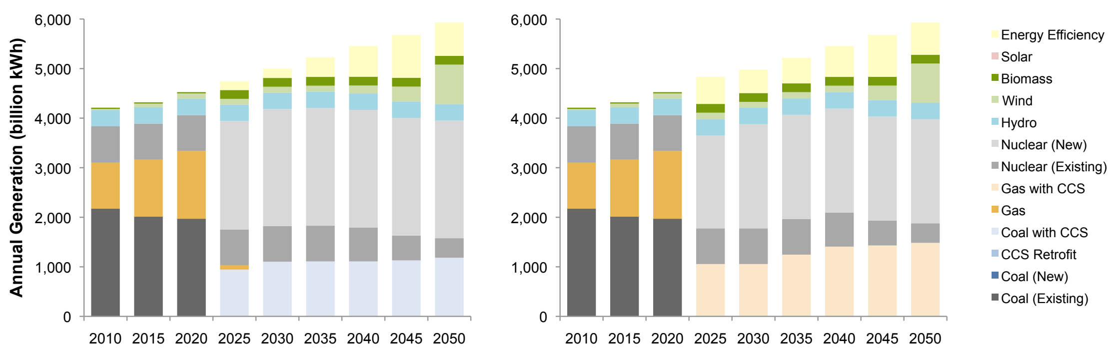

(a) Uncertain upstream emissions (b) No upstream emissions

Figure 4.7: Annual electricity generation (billion kWh) by technology under scenarios with \$40/Mt-CO_2 tax, flat natural gas prices, and mean upstream methane emissions (for the left figure only).

This effect is slightly different from the value of control calculated in previous work (Bistline, 2012), which investigates the utility of upstream CH_4 controls in a policy environment with uncertain cumulative emissions caps.[14] In this situation, the reason that control is so valuable is that, for tight abatement scenarios, this strategy allows existing natural gas plants to generate more during the first stage. It relies

[14]Although cumulative caps may make sense from scientific standpoint, difficulties in implementing such a policy suggest that the actual form of regulation will balance many relevant factors in mechanism design and political feasibility.

on extra capacity from less frequently used units, which currently have low capacity factors and are used primarily as peaking plants, instead of building new ones to keep pace with growing demand. This control scenario has the flexibility of waiting to observe the resolution of uncertainties before building new capacity. It would turn the overbuilding of NGCC units from the mid-1990s onward from a liability into a significant asset for reducing system operating costs, CO_2 emissions (until a more certain policy framework is in place), and conventional air pollutants. This strategy would simultaneously maintain grid reliability without additional capital investments.

In summary, although these system flexibility benefits of capturing CH_4 are not seen within the framework of utilities' decision problem (i.e., where the policy mechanism for internalizing greenhouse gas externalities is a tax), they illustrate how CH_4-reduction or capture technologies are valuable to a variety of stakeholders under a range of policy settings.[15] Reducing upstream CH_4 leakage is valuable under uncertain cumulative emissions constraints due to first-stage flexibility impacts, but under the carbon taxes used here, it is most valuable due to second-stage cost reductions.

4.1.4 Value of Control

The value of control (VOC) is a useful metric for measuring the value of being able to control the outcome of an uncertain situation. It represents the change in value moving from an uncertain state to a desired state without uncertainty. Assuming perfect control, the VOC is determined by comparing the expected cost of the stochastic strategy with all uncertainties and that of the stochastic strategy where the controlled uncertainty assumes a fixed value that minimizes expected cost. The VOC is useful for uncertain parameters that are controllable (either wholly or in part) through allocation decisions. These endogenous uncertainties manifest themselves in the energy-policy domain through technologies that have cost and performance characteristics that can be influenced through directed research and development (R&D) efforts. In this context, the VOC can be interpreted as a proxy for an upper bound on R&D spending.

[15] The importance of CH_4 reductions in an optimal hedging strategy is demonstrated in Labriet, Loulou, and Kanudia (2010), where CH_4 capture at landfills is a robust abatement measure.

Table 4.3: Value of control (billion \$) comparison for selected uncertainties as a fraction of the objective function value (OFV).

Uncertainty	**Units**	**Mean**	**Control**	**VOC**	**% of OFV**
Capital costs (nuclear)	\$/GW	3,980	2,000	297	7.3%
Capital costs (coal with CCS)	\$/GW	3,900	2,000	52	1.3%
Capital costs (solar)	\$/GW	7,800	2,890	39	1.0%
Capital costs (gas with CCS)	\$/GW	2,080	1,000	22	0.5%
NGCC efficiency	% (LHV)	63%	72%	13	0.3%
Acceptance of CO_2 storage	Unitless	>1	1	4	0.1%

Table 4.3 shows the VOC for uncertainties considered in the model. That the nuclear capital cost uncertainty has the highest VOC is not surprising given that it is the most commonly deployed technology during the first stage for both the stochastic and expected-value solutions.

Although likely not as controllable as the other random variables in the model, the climate policy and natural gas price uncertainties have even larger VOCs at \$794 and \$508 billion, respectively. These large values reinforce the importance of these uncertainties for decision-makers and modelers and the high willingness to pay of utilities if these parameters were controllable.

These values beg questions about about optimal R&D portfolio investments instead of upper bounds for R&D for specific technologies. Such questions are addressed with the R&D strategy research in Chapter 6.

4.2 Sensitivity Analyses

The previous section centers on reference model results, which reflect many assumption about the utilities' and generators' capacity planning and dispatch decisions: probability distributions for uncertainties are based on the best-available information; construction costs remain stable over the time horizon; all uncertainties resolve in 2025. To test the robustness of these experimental findings, I conduct sensitivity analyses with respect to important modeling and policy-relevant factors:

1. Using outdated probability distributions

2. Varying the discount rate used in utilities' intertemporal optimization problem
3. Treating construction escalation costs as a separate uncertainty
4. Modifying the availability and technological readiness of nuclear and CCS
5. Varying the resolution date of uncertainties
6. Considering different assumptions for demand growth

4.2.1 Outdated Distributions

The previous sections assume that utilities' probability distributions are based on the most up-to-date understandings and estimates of a range of uncertainties. However, the decision-maker's beliefs may depart from the best-available information for a variety of reasons. Overwhelmed by the number of interrelated factors that must be taken into consideration, utilities may adopt more lax tools for dealing with risks and may instead use fewer uncertainties in the planning process. Resource limitations may lead decision-makers to use heuristic approaches for quantifying distributions, which may entail using values from analyses without sufficiently updated information. Additionally, even if utilities recognize uncertainties and devote substantial resources toward quantifying them, cognitive heuristics and biases can distort probabilities from their "true" distributions, which may impede their ability to accurately quantify and prioritize risk.

In this section, I investigate the impact of using outdated distributions if utilities' beliefs about fuel prices are based on the best-available information from 2007 but the actual, realized distributions are based on the most up-to-date distributions. In other words, first-stage decisions are made based on 2007 expectations for future coal and natural gas prices, but the realizations of random variables in the second stage come from the unexpected distributions discussed earlier, which allows this strategy to be compared with the performance of the optimal stochastic strategy developed in Section 4.1.

All values in this section come from the Energy Information Administration's 2007 Annual Energy Outlook (DOE/EIA, 2007). At that time, energy analysts were bullish on coal and only beginning to understand the long-term impacts of shale gas and how this resource would alter expectations about the domestic energy landscape. According to a National Energy Technology Laboratory report (DOE/NETL, 2007), the resurgence of coal dominated the outlook for the electric power sector, as 145 GW of new coal capacity additions were planned by 2030.[16] According to the 2007 Annual Energy Outlook, the reference (mean) price of coal in 2030 is \$1.87 per MMBtu (in 2010 dollars) with an implicit annual growth rate of -0.9 percent, which is more optimistic than the current mean value of \$2 (0 percent). In contrast, the 2007 Annual Energy Outlook reference natural gas price for 2030 is \$7.07 per MMBtu with a growth rate of 1.7 percent, which is higher and more pessimistic than the current mean value of \$6.10 (1 percent).

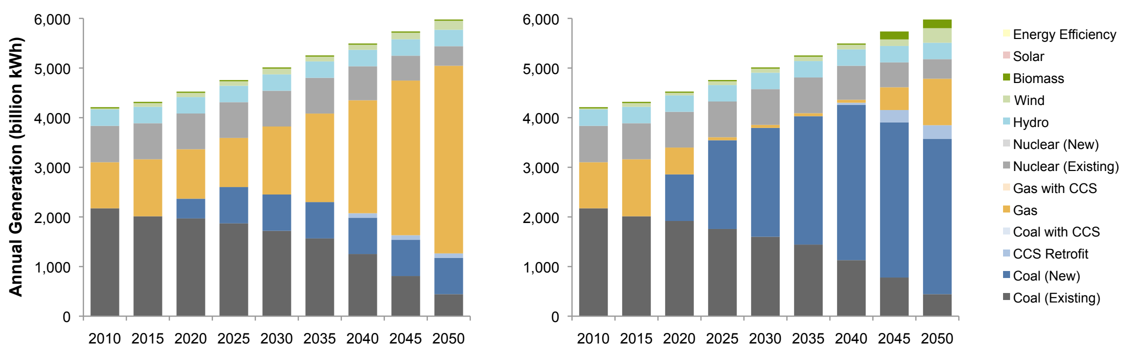

(a) Reference 2013 distributions (b) Distributions based on 2007 expectations

Figure 4.8: Comparison of annual electricity generation (billion kWh) by technology in the no-policy scenario under perfect foresight.

Figure 4.8 shows how this simultaneous optimism about coal and pessimism about natural gas can impact the electricity generation mix in the no-policy scenario. Even these minor differences between expectations of fuel price spreads lead to dramatically different trajectories of capacity additions, generation, and emissions by the end of the time horizon. Under the reference distributions, much of the retired coal capacity (and increasing demand) is replaced by highly efficient NGCC units so that 63 percent

[16]Most of these proposed projects were later cancelled.

of generation by 2050 comes from natural gas. With expectations from 2007, new capacity investments come almost exclusively from supercritical pulverized-coal units so that 64 percent of generation comes from coal by 2050.

Table 4.4: Discounted system costs and VSS (billion \$) comparison under alternate assumptions about probability distributions and the climate policy.

Distributions	**Climate Policy**	z^*	z_d	**VSS**
Reference	Mean value	4,047	4,083	36
Reference	None	4,047	4,108	61
2007	Mean value	4,047	4,127	80
2007	None	4,047	4,289	243

Explicitly incorporating uncertainty in the planning process is especially valuable given outdated 2007 expectations, as shown in Table 4.4. Due to its bullish forecast about coal, the expected-value solution builds more coal with CCS in the first stage with 2007 priors (110 GW) compared with the most up-to-date values (61 GW). As in Section 4.1.1, the VSS comes from avoiding investments in CCS to take advantage of cheap natural gas if available and to avoid stranding these assets in the event that the carbon tax is prohibitively high or low.

This effect is more pronounced when utilities make first-stage decisions assuming that no climate policy will materialize. Under this assumption, utilities would deploy 255 GW of coal capacity without capture equipment instead of nuclear. The costliness of these non-salvagable assets under a range of climate policy scenarios is reflected in the VSS value of \$243 billion using the 2007 distributions.

These experiments indicate an important facet of the existence value of shale gas that is neglected by many analysts—namely, its ability to change expectations about fuel price spreads for the future. As illustrated in Table 4.4, the increased VSS under outdated distributions *vis-à-vis* the reference scenario suggests that shifting expectations between 2007 and the present (i.e., due to new information about shale gas) may have helped to avert costly expenditures that are suboptimal from an expected-value perspective. These VSS values show that, although the explicit incorporation of uncertainty in capacity planning is still important, stochastic planning is not as important now as it was in the early 2000s under a coal-centric investment paradigm. If decision-makers are wrong about their forecasts for future values, the impacts of model error are less severe now than they were 5–10 years ago.

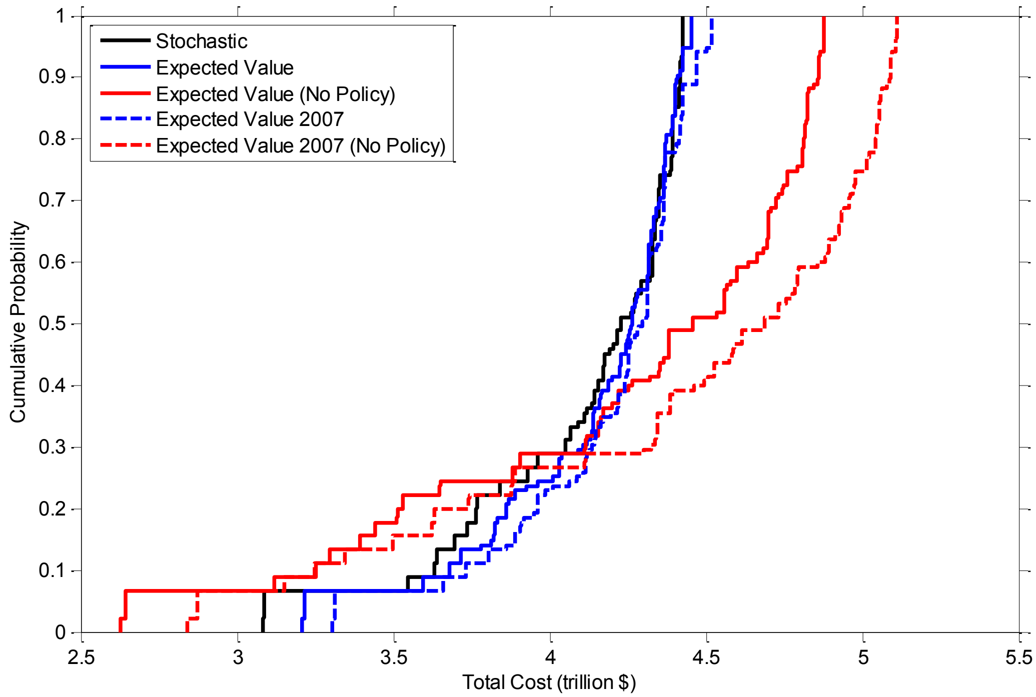

Figure 4.9: Cumulative distribution functions over discounted system costs (trillion $) under the stochastic and expected-value solutions. The solid lines represent the reference distributions, and the dashed lines signify cases where 2007 distributions are assumed when first-stage decisions are made.

Figure 4.9 illustrates the riskiness of adopting stochastic and expected-value strategies under alternate assumptions about probability distributions. Plotting the cumulative distribution functions (CDFs) of the stochastic (black line) and expected-value (blue) strategies serves as an approximate visualization of the VSS, which is the integral between the CDFs of these two approaches.

The stochastic solution not only protects against downside losses but also opens up the possibility of upside gains from volatility. For instance, when natural gas prices are lower than expected, delaying first-stage investments allows the stochastic solution to build more gas-fired capacity, which can take advantage of favorable market conditions. The antifragility (Taleb, 2012) of the stochastic approach provides a strong hedge against uncertainty, reduces risk, and presents simultaneous opportunities to adapt to evolving market conditions. Thus, the stochastic strategy attenuates the adverse effects of downside risk while retaining the option value associated with deferring irreversible commitments until more information is available about potentially lucrative opportunities.[17]

Comparing the stochastic solution with the expected-value solution assuming no climate policy shows the risk premium of the stochastic approach. If the realized values of the carbon and natural gas price uncertainties are low, the stochastic solution will incur investment and operating costs that exceed the expected-value solution by \$0.5 trillion. However, the stochastic solution decreases risk dramatically relative to the expected-value strategies that plan under the assumption of no climate policy or outdated distributions.

4.2.2 Discount Rate

This section considers an experiment in which decision-makers' discount rates are varied from the reference rate of five percent. Table 4.5 corroborates inverse relationships between the discount rate and the EVPI and VSS metrics. The VSS increases more appreciably between five and three percent compared with seven and five percent. The large VSS at a discount rate of three percent is driven by increased investments in coal with CCS under the expected-value strategy, which builds 120 GW in the first stage (compared with 61 GW in the reference case). For lower discount rates, future cash flows have larger impacts on the objective function value and hence on near-term investments, which lowers the incentive to delay investments. Additionally, the VSS and EVPI are larger for smaller discount rates due to objection function value inflation at lower discount rates, as the cost base to which the metrics apply increases with lower rates.

[17] The optionality of capital investments is discussed in Chapter 5.3.

Table 4.5: Discounted system costs (billion $) and comparison of EVPI/VSS metrics for the discount rate sensitivities.

Discount Rate	z_{ws}	z^*	z_d	**EVPI**	**VSS**
3 percent	5,264	5,466	5,576	203	110
5 percent	3,884	4,047	4,083	162	36
7 percent	2,990	3,097	3,114	107	17

4.2.3 Construction Cost Escalation

Raw materials and commodity costs can influence the diffusion of technologies in the electric power sector. Like any input cost uncertainty, construction inputs fluctuate stochastically in time, as policy conditions and the economic environment may be changing simultaneously. IHS CERA Power Capital Cost Index (PCCI) values, which measure project cost inflation for a range of power plants across North America, indicate that construction costs have increased by 124 percent between 2000 and 2012 (IHS CERA, 2013).

Treating construction escalation costs as a separate uncertainty from capital cost uncertainties for individual technologies is useful for a few theoretical and practical reasons. Like stock prices, different types of power plants (i.e., capital assets) have a system cost component and an individual cost component. The construction cost escalation uncertainty incorporates commodity cost volatility like rising steel and copper costs. From a practical perspective, characterizing cost uncertainty through separate individual and systematic components avoids the problem of assessing correlations between cost-related random variables of different plant types by explicitly modeling the cause of the dependence. This separation aids the expert elicitation process in particular, since the elicitation literature suggests that experts are not good at assessing correlational structures (Morgan and Henrion, 1990).

In this experiment, a construction cost index applied to all power plants is varied in two cases:

- *Low construction costs:* Capital costs of power plants decrease linearly from their current levels to reach their 2000 values by 2050.

- *High construction costs:* Capital costs increase linearly through 2050, where they double current levels.

For both the low and high construction cost cases, the VSS increases from $36 billion to $67 billion and $122 billion, respectively. The stochastic strategy in each case has a tendency to delay investment more than the expected-value solution. When construction costs are expected to decline over time, the stochastic strategy can better take advantage of lower capital costs in future years after information about critical uncertainties has been revealed. When construction cost escalations are anticipated, both the stochastic and expected-value solutions significantly decrease first-stage investment. The higher VSS under this case comes from the higher stakes associated with stranded capacity, since down-the-line investments in the second stage are more costly under the higher cost escalation paradigm.

A separate experiment investigates how expectations of higher construction costs interact with a postponed resolution date of 2030. The VSS increases to $231 billion in this case. Delay is most detrimental when the opportunity costs of postponing investments are high. Here, uncertainty generally makes decision-makers want to postpone irrevocably committing resources until future values of random variables are known, but escalating construction costs make utilities want to invest sooner. Delaying the resolution date gives more potential for misaligning investments with true state of the world, and costs are much higher when making investments in replacement capacity, which gives rise to a larger VSS. When uncertainty is not resolved until 2030, the stochastic strategy invests in coal capacity during the first stage. Since the delayed resolution requires some new investments to satisfy rising demand before uncertainty is resolved, the stochastic strategy takes advantage of lower costs in the near-term by constructing cheap capacity with little fuel-cost risk and benefits handsomely under lax climate policy scenarios.

This sensitivity gives a partial response to questions about why utilities would still build coal plants given expectations for unconventional natural gas to depress prices in future years and for future climate policies. Section 4.1.1 also illustrates a case where expectations that no climate policy is on the horizon would build coal units. This section suggests another (less obvious) reason why utilities may find it optimal to build coal plant in the coming years—namely, they may believe that a price on carbon will not materialize in the next couple decades (and perhaps not during the lifetime of the investment) and simultaneously that construction costs will increase.

4.2.4 Limited Technological Availability

The results in Section 4.1 suggest an important role for nuclear technologies to hedge against uncertain climate policies, technological capabilities, and economic conditions. However, many analysts have voiced skepticism regarding the future of nuclear power, suggested that it faces a "crossroads, with possibilities of the start of a renaissance or a slow decline" (NEA, 2012) due to currently low deployment of reactors and to expectations of future cost increases for Generation III/III+ reactors (Anadon et al., 2012).[18] Given these concerns, I consider an experiment in which new nuclear construction is infeasible. The VSS for this sensitivity increases to \$51 billion largely due to the expected-value strategy doubling down on coal with CCS builds (with 108 GW of capacity installations in the first stage), since nuclear units are not available as hedging technologies. The stochastic strategy builds additional wind and relies on increased utilization of existing capacity, which means that this hedging strategy is most valuable when the realized natural gas and carbon prices are low.

A parallel sensitivity analysis considering limited CCS availability suggests that the value of CCS readiness in the second stage is \$13.8 billion. This result illustrates that, although they are not ideally suited for short-term deployment, CCS technologies are an important part of the long-term generation mix.

As discussed in Section 4.1, the reference results demonstrate that utilities have little near-term incentive to build CCS-equipped capacity given uncertainty about climate policy. If learning effects are important for reducing costs to enhance CCS readiness in future decades when greater abatement may be needed, CCS deployment may require near-term public-private partnerships for early pilot and demonstration projects as well as for R&D for capture systems with lower parasitic losses. Rubin et al. (2007) show how historical experience curves for similar technologies like flue-gas

[18]Over the past decade, approximately four reactors per year went online globally (IAEA, 2013).

desulfurization and selective catalytic reduction systems for power plants exhibited cost increases during the initial stages of commercialization. These increases are primarily due to reliability and performance deficiencies in early designs, which are typically not incorporated in long-run learning rates. Thus, capital cost increases for early CCS units (i.e., before cumulative experience with commercial capacity drives down costs) would be consistent with these observations and would be important factors in cost trajectories for CCS.

4.2.5 Uncertainty Resolution Date

In another experiment, the uncertainty resolution period is moved forward and backward from the 2025 reference case to understand the impact of the learning period on the optimal stochastic solution. Due to the structure of the climate policy uncertainty as a tax, advancing the date of uncertainty resolution does not necessarily reduce expected costs, unlike other studies like Labriet, Loulou, and Kanudia (2010).

Table 4.6 shows that changing the resolution date has a small but appreciable difference on the results. The small magnitudes of the changes are due both to the fact that first-stage emissions are not penalized after uncertainty is resolved and to the fact that exogenous retirements do not begin in earnest until about 2030. The VSS increases to $57 billion when the resolution date is pushed to 2030 due to increased investments in coal with CCS (almost 130 GW) under the expected-value strategy. The stochastic strategy delays making new investments as long as possible to obtain information. However, growing demand outstrips existing capacity by 2035 and requires additional investments. Under the 2035 resolution case, the stochastic strategy finally builds coal with CCS, which brings the hedging strategy closer to the expected-value solution and consequently lowering the VSS. Consequently, the VSS is non-monotonic in the resolution date.[19]

[19] Additionally, the EVPI increases when uncertainty is resolved earlier, which mirrors the conclusions of previous research efforts (Parson and Karwat, 2011).

Table 4.6: Discounted system costs (billion $) and comparison of EVPI/VSS metrics for uncertainty the resolution date sensitivities.

Resolution Date	z_{ws}	z^*	z_d	**EVPI**	**VSS**
2020	4,100	4,265	4,301	166	36
2025	3,884	4,047	4,083	162	36
2030	3,778	3,926	3,983	148	57
2035	3,697	3,831	3,857	133	27

4.2.6 Demand

In addition to its seasonal, weekly, and diurnal patterns, electricity demand has a long-term trend of growth. Although the worldwide electricity growth forecast through 2030 is about two percent annually, US growth has slowed from nine percent in 1950s to less than 2.5 percent in 1990s. Between 2000 and 2007, the average US growth dropped to 1.1 percent (Jin et al., 2011).

To test the robustness of the model results to assumptions for demand growth, a sensitivity case uses 2007 projections for demand growth from the Annual Energy Outlook (DOE/EIA, 2007). The forecast of 44 percent growth in total by 2030 is considerably more bullish than the 20 percent increase projected in the 2012 Annual Energy Outlook. Under the high-demand case, the VSS has a small but appreciable increase to $48 billion, and the EVPI increases to $200 billion. Similar increases are also observed when inelastic demand is used with the reference growth case. In each sensitivity, more capacity must be constructed during the first stage to keep pace with growing demand. Increasing demand or decreasing the price responsiveness of demand force utilities to shift irreversible investments earlier in the time horizon, which provides reduced flexibility before more information becomes available and increases the probability that non-salvagable resources will eventually be stranded.

4.2.7 Comparison of Sensitivity Results

The EVPI and VSS results presented in this chapter are summarized in Figure 4.10. The VSS is largest under conditions where the decision-maker does not sufficiently account for the potential for climate constraints in future decades, especially when distributions for fuel prices are outdated and bullish toward coal. Larger values also occur when smaller discount rates (e.g., three percent instead of five percent) are used and when increasing commodity prices put upward pressure on future construction costs. The EVPI is larger when discount rates are low, when demand is higher than projected, and when consumers are less responsive to electricity price changes under policy scenarios (i.e., when demand is inelastic).

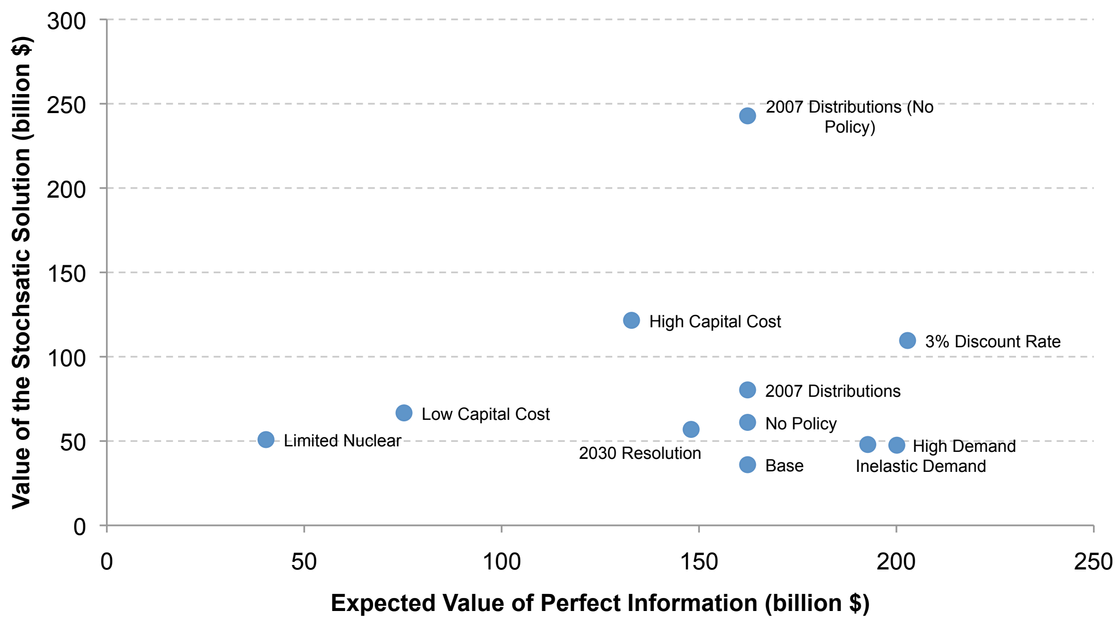

Figure 4.10: Comparison of EVPI and VSS results for the electric sector capacity planning experiments.

Another general conclusion from these results is that uncertainty, particularly in climate policy and technological availability, tends to postpone investments in new generating capacity until more information is made available or uncertainty is resolved. In the reference results, the stochastic strategy builds less capacity during the first stage (431 GW) compared to the expected-value strategy (490 GW), as shown in Figure 4.11. Such reduced first-stage commitments are more pronounced under conditions where the PCCI is high ("High PCCI"), where commodities prices decrease ("Low PCCI"), and where policy constrains investments in new nuclear capacity ("No New Nuclear").

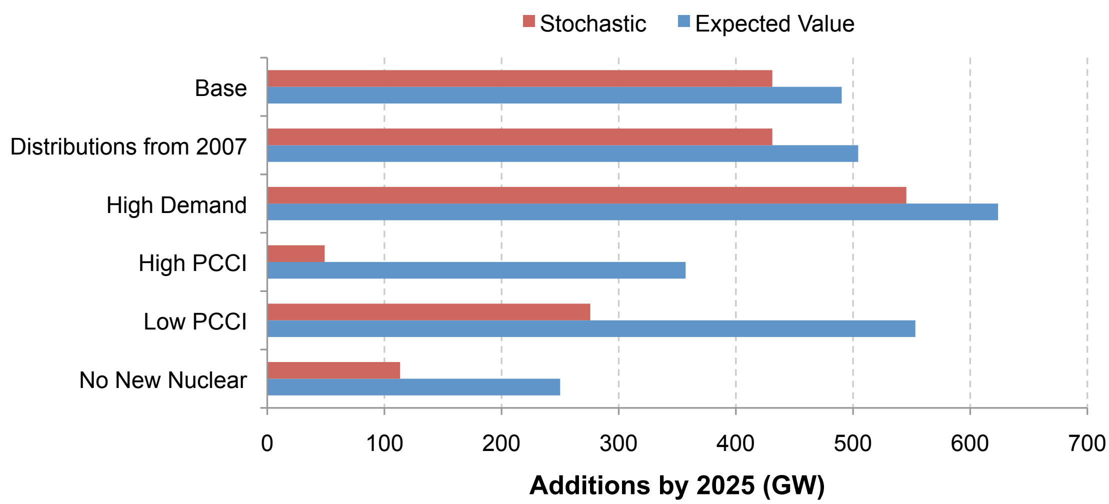

Figure 4.11: Comparison of cumulative capacity additions (GW) by 2025 between the stochastic and expected-value strategies.

4.3 Summary of Findings

This chapter investigates the dynamics of capacity planning and dispatch in the US electric power sector under a range of technological, economic, and policy-related uncertainties. The objective is to determine the sensitivity of near-term decisions to long-term uncertainties by developing stochastic strategies, which account for possible costs of midcourse corrections and hedge against a variety of upside and downside risks. The results suggest important insights about near-term decision-making under uncertainty and the modeling efforts that attempt to inform them.

Using a two-stage stochastic programming approach, model results suggest that the two most critical risks in the near-term planning process are natural gas prices and the stringency of climate policy. Stochastic strategies indicate some near-term hedging from lower-cost wind and nuclear will occur but robustly demonstrate that delaying investment and waiting for more information can be optimal under certain conditions to avoid stranding capital-intensive assets. In particular, the stochastic approach will avoid near-term CCS investments due to the possibility that these assets would be decommissioned either if the climate policy too stringent or too lax or if public opposition prevents cost-effective CO_2 storage. The stochastic strategy instead delays investment in new capacity and relies on increased generation from existing, underutilized (i.e., low capacity factor) natural gas units. One interpretation of the results is that utilities should first pursue quasi-reversible alternatives that provide flexibility and avoid capital-intensive, long-lived investments. Chapter 5 explains these dampening effects of uncertainty in terms of the optionality of investments, leading to more general insights about uncertainty, learning, and irreversibility in the electric power sector.

It is important to note that the model results should not be interpreted as an argument for a do-nothing near-term strategy. As Manne (1996) emphasizes, "Delay should not be confused with inaction." The value of delaying investment to wait for more information applies only to coal with CCS capacity for the stochastic approach *vis-à-vis* the expected-value approach. In fact, both approaches indicate that substantial near-term investments in wind and nuclear are optimal under a robust range of future scenarios. Additionally, the assumption that information will be received in 2025 hinges in part on the existence of sustained R&D efforts in the interim, which is investigated in more detail in Chapter 6.

The largest losses occur when decision-makers' beliefs depart from the best-available information either by using outdated distributions for fuel prices or by adopting optimistic beliefs about the ability to postpone a comprehensive climate policy.[20] These results of misestimation underscore the importance of using actual distributions that incorporate actual data instead of stylized, *ad-hoc* distributions. The VSS comes from the tendency to delay or postpone investments in new generating capacity until more information is made available or uncertainty is resolved. Such hedging policies not only protect against downside losses but also open up the possibility of upside gains from volatility (e.g., when natural gas prices are lower than expected). The stochastic solution is especially valuable if decision-makers do not sufficiently account

[20]These results are especially relevant given the limitations of existing approaches for uncertainty analysis in utility resource planning, as described in Chapter 3.4.

for the potential of climate constraints in future decades, if fuel price projections are outdated, if discount rates are low, or if construction costs are expected to increase over time.

As Chapter 5 discusses in greater detail, these results suggest that a sequential approach to climate policy (e.g., by implementing a new source performance standard in the near future) could incentivize preemptive and supererogatory abatement efforts until more comprehensive climate legislation is in place. These policies may be effective instruments to reduce cost risks for utilities, to safeguard against the erosion of public confidence in political institutions, to demonstrate the feasibility of emissions reductions by beginning with relatively low-cost restrictions, and to lower the probabilities of environmental hazards for society at large.[21]

The model results offer many policy-relevant insights about the future role of unconventional natural gas in the US electric power sector.[22] The value of control for upstream emissions from shale gas is shown to be substantial. Limiting CH_4 leakage allows more natural gas units to be built and operate during the second stage in scenarios where higher carbon taxes are realized and natural gas prices are low to moderate, which means that the development and deployment of these control technologies represent a large value-added proposition for utilities and shale gas developers.[23] Additionally, questions about whether a carbon price will increase or decrease natural gas consumption and whether shale gas availability will influence investments in renewable technologies are shown to hinge on interactions between uncertainties related to natural gas prices and climate policy. The shale gas boom will not impede long-term investments in low-carbon technologies if a sufficiently stringent climate

[21] In US domestic efforts to reduce depletion of the ozone layer leading up to the Montreal Protocol, the initial and persistent focus on banning aerosol applications of chlorofluorocarbons helped to signal commitment toward further remediation efforts. Much like coal applications in the power sector, aerosols represented a substantial source of emissions but had readily available substitutes, which presented a "strong tactical and substantive rationale" for beginning restrictions with this source (Parson, 2003).

[22] These insights are broadly consistent with the results of Huntington (2013) regarding the impact of shale gas on the US electric power sector and on greenhouse gas emissions.

[23] Recent estimates suggest that many strategies for lowering life-cycle CH_4 emissions from natural gas production (e.g., plunger lift systems for liquids unloading, low-bleed pneumatic devices, leak monitoring and repair) are relatively cost-effective with payback periods of less than three years (Bradbury et al., 2013; EPA, 2011b).

policy is enacted in the coming decades. However, if policy-makers fail to provide suitable incentives for firms to internalize climate-related externalities, utilities may overinvest in gas-related infrastructure and underinvest in low-carbon technologies relative to their socially optimal levels.[24] Such effects illustrate the importance of modeling the interactions between multiple uncertainties simultaneously, particularly for complex policy questions.

[24] In addition to adopting a climate policy with appropriate levels of timing, stringency, and credibility, establishing proper incentives requires that non-CO_2 gases be included and also that the global-warming potentials for these gases accurately reflect the latest peer-reviewed research. The 2009 Waxman-Markey bill (United States House of Representatives. 111th Congress. 1st Session, 2009) uses a GWP of 25 for CH_4, which reflects the 100-year timescale value used in the IPCC's Fourth Assessment Report from 2007. The EPA's greenhouse gas emissions inventory (EPA, 2011b) uses a lower value of 21. Chapter 3.4 discusses how these values are smaller than the mean value of 33 from Shindell et al. (2009). The IPCC's Fifth Assessment Report (IPCC, 2013) recommends a value of 34 for the 100-year timescale when feedbacks are taken into account.

Chapter 5

Discussion of Capacity Planning Results

5.1 Uncertainty, Learning, and Irreversibility

The results in Chapter 4 confirm the importance of stock irreversibilities and learning on near-term hedging decisions. If decision-makers eventually learn that climate policies are less severe than initially anticipated *ex ante*, then they will regret costly expenditures on unnecessary control equipment *ex post* and instead will wish that such irreversible investments were delayed. However, if policies are more stringent than expected, decision-makers may regret not taking a more precautionary approach in early periods, as there may be substantial costs associated with delay or with carbon-intensive investments that are suboptimal *ex post*. Thus, irreversibilities in sunk control capital and in investments that impact environmental stocks and flows lead to learning effects that pull in opposite directions—namely, regret over first-stage decisions given updated information and anticipation of second-stage decisions. It is unclear whether these conflicting irreversibilities suggest that near-term investments in abatement capital should be increased (to retain the option to protect against potentially serious impacts from the nondegradable stock of greenhouse gas emissions) or decreased (to wait for more information). Optimal strategies in these contexts must balance prudence and exigency, embodying the spirit of the Latin adage *festina lente* ("make haste slowly").

Given the realizations of uncertainties and near-term decisions, the relative magnitudes of these losses and their associated probabilities will determine the net impact of learning and whether the most appropriate near-term hedging option consists of precautionary investments or delays. These dynamics are compounded by overlapping uncertainties beyond merely climate policy, which may push investment in opposite directions and may be dependent on the incentives of stakeholders.

The model results suggest that the threat of stranded abatement investments outweighs precautionary effects and results in a propensity to delay near-term expenditures. This result mirrors other modeling efforts for power plant investments using other approaches and a more limited number of uncertainties (Knopf et al., 2010; Patiño Echeverri, Fischbeck, and Kriegler, 2009; Tuthill, 2008). However, although investments in coal with carbon capture and storage (CCS) are delayed under the stochastic strategy (relative to the deterministic expected-value case), it is important to note that nuclear and lower-cost wind are still deployed before uncertainty is resolved, which indicates that they are comparatively low-regret hedging technologies. Similar to the results of De Cian and Tavoni (2012), uncertainty about climate policy does not materially impact the first-stage abatement level or generation but mostly affects the portfolio of new capacity additions. The absence of penalties to disincentivize near-term emissions until a price signal is established greatly reduces the salience of the risks of inaction and the impetus for precautionary efforts.

The long-lived nature of electric sector assets and the non-ergodic nature of their evolution suggest that suboptimal investments may become locked in or stranded if assumptions during the planning process are proven incompatible with the realized state of the world, as described by Usher and Strachan (2012). Within the literature on investment under uncertainty in the presence of irreversibilities and externalities (Kolstad, 1996), the welfare losses associated with lock-in and stranded investments are leading justifications for selecting reversible or quasi-reversible abatement options over irreversible ones that are long-lived and capital-intensive. Here, the stochastic hedging solution avoids building CCS-equipped coal capacity in early periods due to concerns that the technology could be stranded if the long-term climate policy is either too high or too low. There is a limited, mid-range band of carbon taxes where CCS can flourish so that, after accounting for uncertainty, it is optimal to wait and see if the realized tax falls in this area before construction. In the meantime, the stochastic strategy relies on increased near-term utilization of existing natural gas combined cycle (NGCC) units and builds only nuclear and lower-cost wind, which are cost-competitive enough to remain online under a wider range of scenarios.

Ultimately, the net impact of learning on either delaying or making precautionary investments depends on the convexity or concavity of the marginal cost function and the shape of the probability distributions for uncertainties (Webster, 2002). Analytical results from the literature on learning in the presence of irreversibilities suggest that concave (convex) marginal costs lead to less (more) of an activity, according to Epstein (1980). However, for a cost-effectiveness framework like this one, where second-stage carbon taxes are uncertain instead of damages from irreversible stock effects (i.e., the distribution over carbon taxes is not equal to the distribution over the social cost of carbon), the marginal costs during the second-stage are not strongly dependent on first-stage emissions. Comparatively low adjustment costs (and the high substitutability of supply-side technologies) allow for a more rapid transition once information is revealed, even if initial installation decisions are shown *ex post* to be incompatible with the realized state of the world.[1]

The direction of the learning effect in this model is influenced primarily by the shape of the probability distribution over the second-stage carbon tax instead of by the curvature of marginal costs. When the expected carbon tax is low and there is a comparatively small probability of a stringent climate policy, then the possibility of learning leads to lower first-stage investments in abatement capital, which is the case here. Due to the concavity of the cost function in the realized carbon tax, the regret from undertaking supererogatory abatement efforts (in which irreversible investments in control capital cannot be recovered) dominates the regret from learning that the carbon tax is higher than the expected value. This occurs because the electric

[1] Although the model captures frictions like construction delay times, it does not account for other constraints like those associated with knowledge stock and learning effects, which may be appreciable if a rapid transition moves toward an energy system with new technologies. The exclusion of these effects likely biases downward first-stage investments and the VSS.

power industry is essentially decarbonized at the mean value of the distribution, which suggests that more stringent taxes would not significantly impact costs for utilities and generators. Notably, this skewness effect in a cost-minimization, policy-compliance setting works in the opposite direction of that in a welfare-maximization, social benefit-cost framework, since the latter incorporates uncertainty with convex loss functions, which have much larger magnitudes of regret in high-damage states, as in Webster (2002).

The small magnitudes of learning effects are due in part to insignificant interperiod interactions in the capacity planning and unit dispatch model. Although the bottom-up detail in the representation of capital stock permits more interperiod effects relative to top-down models, there are still many important processes that are excluded from the model, which would cause the first-stage strategy to have a greater impact on second-stage marginal costs. First, vintaging effects and infrastructural inertia of the capital stock, particularly if constraints lead a portion of the residual capital stock to become less malleable, may delay shifts in technological adoption when prices change, which would lead to higher marginal cost reductions in later periods. Even though the capacity planning model has a high degree of flexibility (as discussed in Section 5.3), it has less malleability and more real-world frictions (e.g., construction lead times, upper bounds on capacity expansion) than many top-down models due to factors like explicit modeling of retirements, construction lead times, and constraints on building new capacity. Second, the model does not include damages from climate change, which can lead to dramatic interperiod interactions if the model represents a threshold response. Finally, the rate of technological improvement or energy efficiency adoption may be influenced by climate policies. The incorporation of endogenous technological change, for instance, can lead to more precautionary investments in the first stage (Webster, 2002). Thus, the conclusion to be drawn from this work is not that the likelihood of benefitting from stochastic hedging is small for near-term decisions but that the current generation of energy-economic models (even those that have a fair amount of detail in capacity planning and unit dispatch decisions) may be inadequate for capturing the details that are most relevant to questions of uncertainty and sequential learning.

5.2 Inducing Anticipatory Actions

Uncertainties that lead to first-stage precautionary efforts can be conceptualized as ones that induce so-called "anticipatory actions" (Labriet, Loulou, and Kanudia, 2010). Random events that induce anticipatory actions (e.g., a climate policy uncertainty that causes supererogatory abatement relative to the deterministic case) merit explicit treatment and modeling of uncertainty in a sequential manner. In this sense, uncertainties with appreciable VSS metrics can be linked to ones that cause anticipatory actions. As the previous section emphasizes, the decision about whether to take action now or to wait for better information depends on the risk and opportunity cost of delay. In the case of climate policy uncertainty, the model results demonstrate an anticipatory absence of investment rather than anticipatory action, which leads to a nonzero VSS.

These conclusions underscore and reframe the notion that there are no silver bullet abatement technologies in the electric power sector. Evocations of this no-silver-bullet mantra typically emphasize the notable dearth of technologies that have the scale, wide applicability, cost characteristics, and technological readiness to reduce emissions dramatically in the context of meeting a deterministic emissions target. In this paper, the lack of silver bullet technologies is reflected in the conspicuous absence of near-term hedges that are robust for all uncertainties. Nuclear and low-cost wind are the only technologies deployed widely under the stochastic hedging strategy, though many large-scale nuclear builds may be costly in the next several years due to the increasing competitiveness of electricity markets and utilities' small market capitalizations. Also, model results suggest that technological cost and performance uncertainties do not induce anticipatory actions, with each having a zero-valued VSS. This is true when examining uncertainties within a cost-effectiveness framework (i.e., since capital cost uncertainties for future periods do not alter near-term abatement) or when there is no construction escalation over time.

5.3 Optionality of Capital Investments

Why is the value of information gathering (i.e., the expected value of perfect information or EVPI) so much greater than the value of including uncertainty (i.e., the value of the stochastic solution or VSS) in electric sector planning? As shown in Figure 5.1, analyzing the EVPI and VSS involves comparing the expected costs of the expected-value (z_d), stochastic (z^*), and perfect information (z_{ws}) strategies.

The EVPI essentially measures the opportunity cost of delaying action. If access to early information equips utilities to make different decisions, the information has value. Otherwise, the EVPI is zero, and the decision-maker can delay decisions costlessly until uncertainties are resolved. The high EVPI found in this analysis suggests that there is a limited availability and adequacy of the hedging options in the electric power sector.[2]

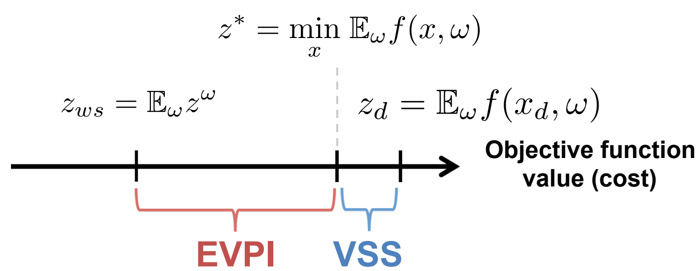

Figure 5.1: Number line comparing values that comprise the EVPI and VSS metrics. The spacing between values is illustrative.

The VSS measures the degree of asymmetry in the decision-maker's loss function and the degree to which suboptimal decisions impact the objective function value. The VSS is large (i.e., the expected-value solution is a bad approximation for the stochastic solution) when:

1. The optimization problem exhibits nonlinear or nonconvex behavior

[2]This is another way of expressing the notion that there are no alternatives in the electricity generation choice set that are robust across all of the uncertainties considered here. There is no single technology that is perfectly suited to all states of the world.

2. Probability distributions on uncertain quantities are asymmetric (i.e., exhibit considerable skewness)

3. Dependence exists between random variables

4. Random variables have large supports

The VSS is small here due to many factors. Regarding point (1), the capacity planning problem is predominantly linear in nature, since many of the nonlinearities in the climate system and utility functions are not included from the decision-maker's perspective, which is framed within a cost-effectiveness framework with uncertain carbon taxes. As described in Chapter 3.4, the distributions used in this analysis are largely symmetric, which means that (2) does not apply. Also, the distributions were created to avoid any explicit dependence between random variables (3) and to bypass large supports (4).

It is important to note that the EVPI and VSS are a small fraction of the objective function value. The EVPI of $162 billion is 3.8 percent of the objective function value of $4.3 trillion, and the VSS of $36 billion is 0.8 percent. Figure 5.2 shows that 64 percent of the total costs come through dispatch and maintenance costs of capital (including fuel costs), whereas capital investments comprise just 24 percent. As a result, if first-stage decisions are suboptimal for the realized scenario *ex post*, the total losses amount to the stranded costs of assets, which are only a small fraction of total costs across the time horizon. For instance, if the decision-maker builds too many NGCC units only to learn that gas prices are higher than expected, then the financial losses for a single power plant would only amount to perhaps $1 billion of the $4,300 billion objective function value. Recourse decisions that are made after information is revealed allow the system to adapt and to avoid incurring increased operating costs in perpetuity if initial decisions are wrong.[3]

[3]The decision to shut down a facility can be viewed as an investment in the sense that initial payments to end contractual commitments generate prospective utility in the form of decreased future losses (Dixit and Pindyck, 1994). In the processes of capital budgeting and resource planning, the options to abandon an investment or halt operation as market conditions deteriorate (e.g., if factor costs like fuel prices exhibit sustained growth) can be valuable for avoiding fixed and operating costs while potentially gaining revenue from the resale value of capital equipment.

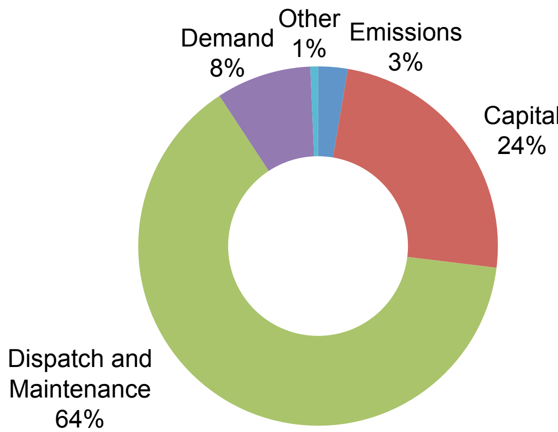

Figure 5.2: Decomposition of the $4.3 trillion objective function value of total discounted costs.

This effect suggests that there are two relevant option values associated with power plant investments. The most commonly discussed and modeled value is the option of firms to invest in capital-intensive and essentially irreversible generators that can be delayed (Tuthill, 2008), which is analogous to a financial call option. This perpetual call option gives utilities the right, but not the obligation, to pay a specified amount (i.e., the strike price of the overnight capital cost) to receive an asset (i.e., a power plant) with uncertain future cash flows due to stochastic processes like prices of emissions permits and fuels. If and when the firm exercises this option to build the plant (i.e., when the asset's value sufficiently exceeds the exercise price and the option is "in the money"), the firm gives up the opportunity to wait for additional information about the future values of unknown quantities.[4]

The above analysis suggests that the second relevant concept is that of a put option. After constructing a power plant, generators have the perpetual, costless, and quasi-reversible put option to generate electricity. They have the right, which they need not exercise, to pay the strike price (i.e., fuel costs) to receive an asset (i.e., revenue from generating and selling electricity). This perpetual put option allows generators to refuse to generate at a loss if the investment is later revealed to be

[4]The claim that regulatory uncertainty depresses investment by creating an option value of delay is investigated empirically by Ishii and Yan (2004), which demonstrates this effect for investments between 1996 and 2000 in the United States under uncertainty about regulatory restructuring.

incompatible with the realized state of the world (e.g., if fuel prices or carbon taxes are too high).

The call option pertains to the initial capital investment decision, while the put option is relevant to the operational decision by firms once the unit has been constructed. Both options must be accounted for in the capacity planning problem. However, the choice of when and if to invest in the call option is essentially the capacity planning problem, which implicitly incorporates the downstream put-option (recourse) decision.

Ultimately, the resiliency that makes such options possible in electric sector planning comes through a variety of industry-specific sources:

- The ability to react and make decisions after new information becomes available
- Low opportunity costs associated with irreversible investments
- Quick construction lead times relative to operating lifetimes

5.4 Policy Implications

The results in Chapter 4 can help to inform policy instrument choices and timing decisions associated with climate change. Given dynamics that discourage precautionary capital investments in control equipment in the presence of uncertainty, these results suggest that a price on carbon may not be enough to properly incentivize utilities to internalize greenhouse gas externalities. Figure 4.9 illustrates that the most substantial cost risks and expected losses occur when utilities and generators make near-term decisions believing that the prospects of a low (or nonexistent) climate policy are increasingly likely.

Given that this outcome would likely be socially suboptimal, a tiered approach to climate policy may be a more effective means of meeting many simultaneous goals. A layered (i.e., tiered) approach would offer simultaneous incentives and policy mechanisms for reducing emissions at federal, state, and local levels like carbon taxes, research and development (R&D) subsidies, and technological standards. Although the policy redundancy may make any given policy less efficient, the potential co-benefits outside of reduced environmental damages may outweigh such deadweight losses (e.g., subsidized R&D can remedy innovation externalities).

A second policy implication for the electric power sector is that sequential or staged approaches may be useful in reducing greenhouse gas emissions even before comprehensive legislation has been passed to put a price on carbon. The policy goal of reducing carbon emissions may be more politically feasible and may induce greater compliance if multiple policies are staged sequentially over time. Again, this redundant and precautionary approach to avoiding irreversible investments in carbon-intensive capital could be important if more ambitious efforts to curb emissions through direct emissions-pricing policies prove to be politically infeasible in the future. Uncertainty about political processes, for instance, may justify these complementary regulations. For example, before enacting climate legislation, the United States Environmental Protection Agency has proposed new source performance standards under Section 111(b) of the Clean Air Act, which would effectively prevent new additions of coal-fired power plants without CCS. If enacted soon, these standards for new power plants would prevent one of the largest carbon-emitting technologies from retaining or expanding its market share even if a long-term climate policy fails.[5] Performance standards have the benefits of enhancing institutional credibility for establishing firm commitments to reducing emissions and also of ameliorating the potential problem of firms racing to install capacity before more restrictive rules come into effect (e.g., if firms anticipate grandfathering or exemptions). Again, although an economic analysis of seemingly redundant policies in a first-best setting would suggest that this approach imposes additional costs on firms and consumers, this analysis shows how such policies can be economically efficient in a second-best setting where the actual timing, stringency, or implementation of a carbon tax are suboptimal.

[5] New source performance standards also can be effective, because they leave open the possibility of constructing CCS-equipped coal instead of banning coal outright, which may make such approaches more politically viable and economically efficient. Lashof et al. (2012) detail the legal basis for setting such power plant regulations and also discuss how standards can provide an incentive to invest in precautionary abatement capacity due to banking provisions. Another regulatory action that would achieve similar results would be the proposed Mercury and Air Toxics Standards for Power Plants.

The use of nested policy instruments with periodic reassessments is consistent with an adaptive approach to managing climate risk. The objective of adaptive management is to design policies, decision architectures, and institutions that encourage monitoring and active learning about systems and allow for flexibility, adjustment, and adaptation as new information becomes available. Sequential decision-making frameworks like the one used here are ideally suited to provide insights and decision support for adaptive management. The ozone regime established through the Montreal Protocol is a noteworthy example of an effective adaptive management framework and the only major international environmental effort to date to adopt such a system with repeated negotiations and dynamic ratchets that adjusted controls, incorporated additional chemicals, and developed new institutions and mandates over time (Parson, 2003). Adaptive management also improves upon some of the shortcomings of a contingent agreement approach to dynamic policies in environments of extreme complexity and uncertainty. Approaches based on contingent responses require the *ex-ante* specification and enumeration of all possible scenarios and optimal responses many years in advance, which neglect the potential for unexpected sources of information and the possibility that such information may generate greater uncertainty and novel questions.

Chapter 6

Energy Technology R&D Portfolio Management

This chapter introduces a stochastic research and development (R&D) portfolio management framework and presents results for energy technology R&D strategy in a carbon-constrained world. In previous chapters, the capacity planning model made first-period investment and operational decisions under uncertainties about future policies, technologies, and fuel prices. Technology development could lower costs when second-stage decisions are made. However, the model assumed that first-stage decisions did not influence the probability distributions over potential outcomes, which meant that future technological characteristics were insensitive to near-term decisions. Figure 6.1 illustrates the decision diagram for the capacity planning model developed in Chapter 3, where technological and economic uncertainties were exogenous.[1]

[1] Boxes in the diagram represent decisions and are assumed to be part of a no-forgetting network (i.e., any information available to parents of a decision node are also available to the node itself). Uncertainties are represented as ovals, though the simplified representation in Figure 6.1 groups all uncertainties as either technological or economic. The hexagonal nodes are special deterministic nodes representing the values that decision-makers are optimizing.

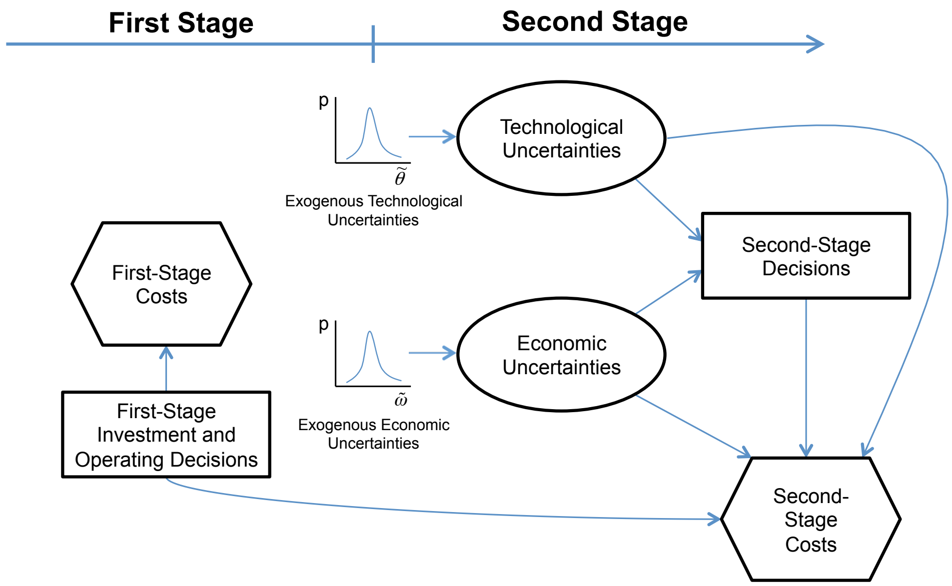

Figure 6.1: Decision diagram for the capacity planning model from Chapter 3.

This chapter describes and uses a model in which investments in R&D programs can be made at the beginning of the time horizon.[2] The decision-maker then sequentially learns whether the R&D program is successful by 2015 when making first-stage capacity installation decisions. The resolutions of all endogenous technological uncertainties and exogenous market uncertainties occur before making second-stage decisions. Figure 6.2 shows the decision diagram for this model.[3] This chapter formulates the R&D decision model labeled "R&D Investment Decisions" in Figure 6.2. The R&D allocation model uses empirically-derived innovation production functions to compare the benefits and costs of energy technology development programs.

[2]The R&D strategist is assumed to make a single R&D allocation decision. Blanford and Weyant (2007) use a sequential decision framework with two R&D decision periods but find the impact of future R&D choices on near-term investments to be small. This result suggests that the single-stage simplification used here is reasonable for tractability while maintaining important dynamics of the R&D decision.

[3]Note the influence arrow on the diagram from the R&D portfolio allocation decision to the technological uncertainties. This arrow means that the initial R&D decisions influence probability distributions for the technical cost and performance uncertainties.

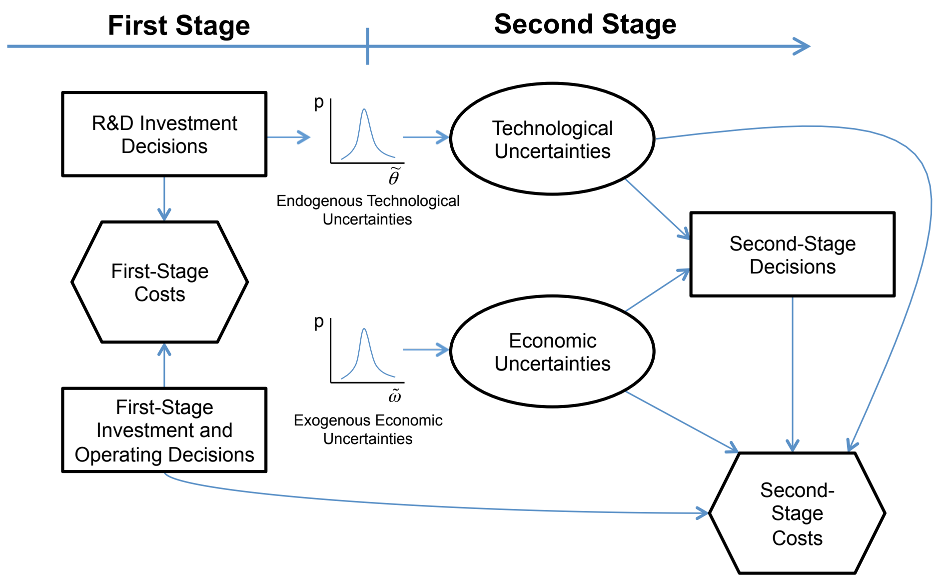

Figure 6.2: Decision diagram for the integrated R&D strategy and capacity planning model from Chapter 6.

The model is distinguished from frameworks by other researchers in its stochastic diffusion mapping through the two-stage capacity planning model described in Chapter 3. The key attribute of the model is that, when R&D funding and first-stage decisions are made, the realizations of other exogenous uncertainties (e.g., abatement stringency) are unknown, which translates into uncertainty about diffusion markets for the technologies upon which R&D acts.

Section 6.1 situates this work within the context of the existing energy R&D literature and highlights the contributions of the framework and results. In particular, this section describes the advantages of a stochastic representation of market diffusion and R&D success valuation. Section 6.2 provides a mathematical formulation for the R&D portfolio model and determines first-order conditions for optimality. Section 6.3 discusses the results of these modeling efforts and their policy implications. Finally, a summary of these findings and potential extensions are briefly described in Section 6.4.

6.1 Introduction

6.1.1 Background

Managing technological change is important for industry, government, and society. In addition to expanding an economy's production-possibility frontier and encouraging growth, the development of new technologies and improvement of existing ones can enhance policy responses to climate change. The current set of technological alternatives is likely insufficient for achieving meaningful abatement to manage climate risk at a socially acceptable cost. However, research, development, and demonstration efforts can introduce new options in future periods. Abatement and technological R&D are inextricably linked in that a given level of abatement influences the set of technologies available for deployment, while the given set of available technological options simultaneously influences the optimal level of mitigation.

No matter if an environmental externality is fully or partially internalized, environmental policy alone (e.g., a carbon tax) likely does not provide sufficient incentives to foster innovation and technological diffusion at a socially optimal level.[4] The effect of a contemporaneous carbon price signal on R&D by private firms is insufficient to induce longer-term invention and innovation activities, especially for basic R&D.[5] Innovation market failures require complementary instruments that typically have a normative rationale of correcting the misallocation of private resources to R&D. The most commonly cited innovation market failures and barriers leading to suboptimal

[4] The literature on induced innovation suggests that alternative policy instruments have markedly different impacts on the incentives for innovation and diffusion of "clean" technologies (Fischer, Parry, and Pizer, 2003; Jung, Krutilla, and Boyd, 1996; Milliman and Price, 1989). For instance, using empirical evidence from cap-and-trade programs for sulfur dioxide and nitrogen oxide, Taylor (2012) suggests that these policies may not provide sustained incentives to encourage R&D and may induce additional uncertainty.

[5] Conversely, technology-oriented policies like low-carbon technology R&D subsidies may lead to emissions reductions in the absence of a climate policy but are exceedingly inefficient in achieving abatement goals if used as substitutes instead of complements. The combined implementation of an emissions price and technology policy appreciably lowers costs compared with the emissions-price instrument alone (Fischer and Newell, 2008).

R&D allocation levels are appropriability, technological lock-in, knowledge externalities, adoption externalities, and incomplete information.[6]

Despite the considerable research attention dedicated to optimal instrument choice for internalizing emissions-related externalities under uncertainty (Baker, 2009), there is comparably little research on strategies for energy technology R&D portfolio management, even though the future technological state is an important factor in policy choice alongside the discount rate, uncertainty, and the assumed benefits of abatement. For large-scale energy-economic and integrated assessment models, many frameworks assume that technological cost and performance characteristics improve exogenously over time (i.e., the autonomous rate of technological change is not influenced by changes in policy or changes in relative prices) or that endogenous technical learning will lead to technological change with increasing deployment. It is uncommon to link R&D decision models with energy-economic models.

Most papers that examine R&D decisions under technological and policy-related uncertainty are theoretical and only examine allocations for a single technology, as discussed in the review by Baker and Shittu (2008). Baker and Solak (2011) summarize the effect of uncertainty on energy technology R&D portfolios. Blanford (2009) and Blanford and Weyant (2007) formulate an R&D decision framework but use illustrative parameters and assume exogenous market uncertainties are resolved when R&D decisions are made. Baker and Solak (2011) use elicited data in a stochastic R&D decision model but do not consider exogenous market uncertainties. Baker and Solak (2013) use a stochastic programming version of DICE to examine R&D portfolio management. However, the only exogenous uncertainty considered is climate change damages, and the sequential decision-making model represents technologies in a highly aggregated manner with marginal abatement cost curves.

[6] These theoretical arguments supporting the proposition that private firms underinvest in knowledge creation and innovative activity are confirmed by empirical research, which suggests that social rates of return to R&D investments exceed private rates due to third-party spillovers (Bloom, Schankerman, and Van Reenen, 2013; Griliches, 1992; Mansfield et al., 1977). Although it is challenging to quantify the relative significance of these market failures and consequently to prioritize the effectiveness of potential interventions, it is at least possible to link policy instruments to specific market failures (Jaffe, 2012). The allocations toward public R&D discussed in this chapter most closely target issues of appropriability and capital market failures.

6.1.2 Motivations and Contributions

The primary objective of this research is to use the framework developed in previous chapters to inform questions of energy technology R&D strategy. Uncertainty is a fundamental characteristic of the R&D process. The stochastic and dynamic aspects of these questions are significant structural features of R&D strategy, including uncertainty in market and policy conditions, the relationship between R&D investments and technological outcomes, and the ability to adjust decisions over time based on learning. Despite the centrality and policy relevance of uncertainty, there is a need for new tools to cope with uncertainty explicitly and to provide decision-making support for R&D investments (National Research Council, 2007, 2005).

This research also informs outstanding questions about how to value technological advances. Measuring the benefits of R&D expenditures typically involves two distinct steps—namely, modeling the relationship between R&D portfolio investments and their potential outcomes as well as valuing these outcomes. This research investigates R&D success valuations in a sequential decision-making setting. This approach provides a more accurate representation of the R&D decision-maker's dilemma in which allocation decisions must be made in an uncertain market environment, where prospective conditions are subject to many contemporaneous sources of uncertainty.[7] Section 6.1.4 discusses the benefits and novelty of this approach.

Another objective is to parameterize innovation production functions using results derived from expert elicitations rather than using *ad-hoc* values. Previous innovation production function analyses (Blanford, 2009, 2006) use stylized values that are the same across all technologies. Although these analyses illustrate the framework, they offer limited insight into actual R&D allocation decisions. In contrast, the work here provides some empirical grounding for the chosen values that link model representations with on-the-ground expectations for R&D program characteristics. This trait, combined with the stochastic valuation model and larger portfolio of R&D programs,

[7] A related contribution is to develop a better understanding of how regulatory uncertainty influences innovation. Recent work has shown how demand-pull actions from regulation can be as large of a driver of innovation as technology-push R&D expenditures (Taylor, 2005), which suggests that accounting for the relationship between regulatory uncertainty and R&D activity is important.

suggests that model outputs offer a greater degree of normative decision support compared with previous analyses.

An overarching goal of this chapter is to embed the treatment of R&D control within the broader context of energy modeling. There is a long history of using energy-economic and integrated assessment models to quantify the benefits of technological developments, but only a limited amount of work that explicitly links these models with an R&D portfolio framework. The research here leverages the experience, tools, and insights from the energy modeling community to address the R&D portfolio allocation problem.

6.1.3 Conceptualization of R&D Success

There are many ways to conceptualize the success of an R&D program (Chan et al., 2011). First, R&D success can be modeled as increasing the (binary) probability of success in achieving specific technical or cost metrics for specific technologies. For instance, Baker, Chon, and Keisler (2009a) define R&D success for advanced solar technologies as meeting fixed targets for efficiency, operating lifetime, and manufacturing cost. Second, success can reduce the number of years required to reach technical or cost targets, which formalizes the notion that R&D success does not provide benefits in perpetuity. Blanford (2009) adopts this framework in characterizing one "optimistic" technological pathway (i.e., with successful R&D) and another "pessimistic" pathway (i.e., one that achieves the same targets with a delay). A final way to conceptualize R&D success is as adjusting distributions over cost and performance metrics, which is the definition used in this work. This stochastic conceptualization has not been implemented as widely as the others; however, some recent expert elicitations have adopted this definition (Anadon et al., 2011), which allows these elicited distributions to be used in this framework.

There are many benefits to this conceptualization of R&D success. First, many scientific and engineering experts have intuitions for ranges of possible outcomes from research endeavors but have a more difficult time assessing probabilities of reaching *a priori* cost and performance targets (Chan et al., 2011). This tendency suggests that elicitations would be facilitated and the quality of their outputs enhanced by structuring these probability assessments according to this conceptualization. Second, the first two conceptualizations of R&D success mentioned earlier are subsumed by this probabilistic approach. Endogenous probabilities of reaching fixed targets or accelerated development can both be incorporated into this flexible framework. Finally, the largest advantage of this probabilistic framework is its versatility in describing the impacts of R&D success. As Figure 6.3 suggests, this conceptualization is capable of representing a diverse set of effects on distributions:

- Shifting the mean (or other measure of central tendency) of a distribution, as shown in the cost reduction from C to C' in Figure 6.3
- Reducing the variance
- Eliminating fat tails (i.e., removing the possibility that a technology is always too expensive for deployment), which is shown in the elimination of probability mass from the far-right side of the distribution in Figure 6.3
- Combining the aforementioned effects

This versatility suggests why this approach, while more intuitive and capable of complex representations, has not been implemented widely in a modeling setting. This conceptualization is most useful when the probabilistic information can be fully incorporated in the planning process, which requires a framework where uncertainty is treated explicitly (e.g., the stochastic programming setting here).

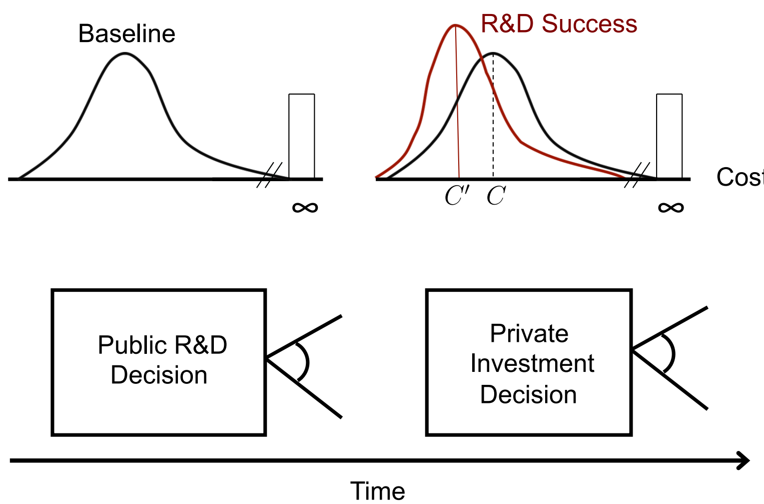

Figure 6.3: Illustrative example of the effect of an R&D success on an energy technology's cost distribution.

Conditional on the success of an R&D program (which is determined through the R&D portfolio model described in Section 6.2), the valuation model uses the R&D success distribution. If the program fails, the R&D expenditure is considered to be a sunk cost, and planning decisions are made with the technological baseline distribution. Realistically, this relationship is more complicated, as there are many levels of success and failure with corresponding distributions (i.e., a distribution of distributions that is conditional on the R&D allocation). However, this more complicated treatment would be difficult to implement due to the large number of probability assessments and also to the computational intensity of the resulting problem.

6.1.4 Assessing Benefits and Modeling Market Diffusion

In the literature on technology R&D optionality, Dixit and Pindyck (1994) note that there are two uncertain underlying assets involved. The first is the state of technology, which R&D expenditures influence. The second is the set of exogenous market conditions (e.g., advances in substitute/complementary/enabling technologies, public acceptance, and regulatory environment). The energy technology R&D literature has primarily focused on the first uncertainty and treated the second deterministically using existing energy-economic models. However, nearly all energy-economic and integrated assessment models do not represent the stochasticity or hedging potential of technologies and do not treat market uncertainties explicitly through sequential decision-making frameworks (Kann and Weyant, 2000). This work is the first to incorporate uncertainty explicitly into a technology-rich R&D valuation model.

Market acceptance is a key determinant of the economic success and valuation of R&D programs. Uncertain economic and political conditions in the future market landscape are important for assessing the diffusion potential of energy technologies. Whether in *ex-ante* prospective R&D decisions or in *ex-post* program evaluations, all applicable economic sectors, regions, and technologies and their associated uncertainties should be concurrently analyzed to capture the diffusion potential of technologies.[8] Diffusion depends on many interrelated factors that may be exogenous to R&D itself, including parallel developments in related technologies, factor costs, the policy environment, demand characteristics, and more general macroeconomic trends. For example, the stringency and timing of an internalization policy for greenhouse gas emissions is an inherently uncertain but important determinant of the value of R&D success. Taylor (2005) provides a summary of the literature linking regulatory stringency, expectation, and uncertainty with innovation.

In a report requested by Congress to develop a methodology for assessing prospective benefits of energy R&D, the National Research Council (2005) refers to the processes of characterizing and incorporating uncertainty as "essential features of prospective benefits evaluation." The report highlights three types of uncertainty: uncertainty about the outcome of a specific R&D program, uncertainty about a technology's market acceptance, and uncertainty about the future state of the world, which incorporates factors that are unrelated to the technology itself. Given the importance of market risks, the research in this chapter is the first to implement these

[8] Measuring the benefits of R&D more generally is challenging to conceptualize and quantify (National Research Council, 2005). In particular, spillovers and social benefits require a detailed understanding of complex linkages in an economy and the cascading secondary impacts of technological advance, which makes quantification more difficult as evaluations move away from the innovation itself. The measurement problem of connecting research in the present with uncertain market conditions in the future is inherently difficult for private firms and perhaps even more challenging for public decision-makers who typically have lower access to information about market conditions.

modeling suggestions by incorporating all three sources of uncertainty in a unified modeling framework.[9] This research also extends the National Research Council's recommendations by offering a portfolio design model to consider R&D allocation questions for a range of technologies simultaneously, which more accurately captures the interactions between technological uncertainties and the corresponding changes in expected benefits of R&D.

Another motivation for carefully quantifying and assessing uncertainty is that uncertainty about market conditions may be a barrier that causes underinvestment in R&D by private firms (Cohen and Noll, 1991).[10] Like all ventures, R&D investments face many risks that lead investors to demand a risk premium, which requires an estimation of risk by forecasting market conditions in which diffusion may occur. These uncertainties are especially challenging to quantify, since there is little relevant historical data to construct analogues for many inventions and innovations. Thus, public R&D projects are likely exposed to greater risk and require analysis that can explicitly aid in decision-making under uncertainty.

Given that the valuation of an R&D program depends on its diffusion potential, models that assess prospective deployment rely on assumptions about the decision-making approaches of economic agents and about the treatment of uncertainty in the planning process, as illustrated in the capacity planning results in Chapter 4. The explicit inclusion of uncertainty using a here-and-now approach means that exogenous market uncertainties like the climate policy are unknown when R&D allocation or first-stage decisions are made, as shown in Figure 6.4.[11] This approach closely reflects the situation faced by many decision-makers (i.e., since the future state of the world

[9]The National Research Council (2005) states that, "Market risk factors are often critical to evaluating the potential of an R&D program. Indeed, for investments in fairly specific technologies, the risks associated with market acceptance may overwhelm those associated with technical success."

[10]Information asymmetry about a technology's potential introduces issues of adverse selection and the winner's curse, which may raise the cost of capital for financing the development of new technologies (Jaffe, 2012).

[11]Since deployment and dispatch decisions are made by utilities and generators, this framework assumes that realizations of market uncertainties cannot be influenced by the decision-maker. Although these uncertainties are treated exogenously here, climate policy choice is endogenous to the R&D strategy problem from a social planner's perspective, as optimal abatement is influenced by the state of technology and expectations about future developments (Blanford, 2006).

cannot be accurately predicted *ex ante*) and is used in this modeling work. In contrast, the wait-and-see approach to uncertainty analysis assumes perfect information about economic and policy uncertainties when initial decisions are made. This approach is implicitly assumed in models like Anadon et al. (2011) and Blanford (2009).

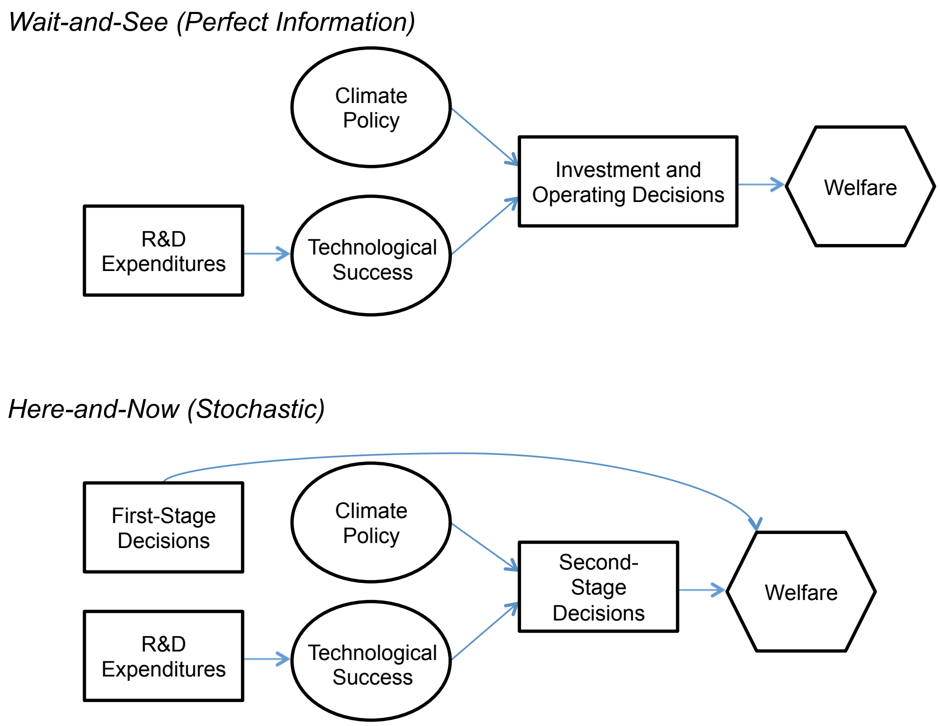

Figure 6.4: Decision diagram comparison of wait-and-see (i.e., perfect information) and here-and-now (i.e., stochastic) R&D valuation models. Uncertainty about the stringency of climate policy is used as an example of a market risk that is exogenous to technological uncertainties but is a key determinant of diffusion.

It is unclear *prima facie* whether the explicit inclusion of uncertainty using a stochastic hedging approach increases or decreases the expected value of R&D success. As the metrics in Chapter 3.3 suggest, it is unambiguous that the failure to consider uncertainty in the decision-making process leads to suboptimal performance in expectation. However, different decision-making approaches may understate or overstate the value of market penetration if uncertainties and competition are not considered. The breeder reactor program in Section 6.1.5 is an example where insufficiently accounting for diffusion risks caused a systematic upward bias in the valuation of technological change, which ultimately translated into higher than optimal R&D funding. This conclusion is not the case generally, as the proof in Appendix C demonstrates. The influence of different decision-making approaches on the value of R&D success and ultimately on optimal R&D investment depends on:

1. How uncertainties interact (i.e., the objective function, constraints, and parameterization of the optimization problem)

2. Form of the distributions chosen

3. Change in technological characteristics brought about through R&D programs

Ultimately, this ambiguity is a significant motivation for the modeling research presented in Section 6.3 to understand the link between the value of R&D success and decision-making approaches under uncertainty.

6.1.5 Case Study of the Breeder Reactor Program

The United States (US) breeder reactor program illustrates a federal R&D program in which overconfidence and the absence of careful consideration of exogenous market uncertainties created suboptimal investments.[12] The diffusion of the breeder reactor and the value of associated R&D efforts hinged on factors related to electricity demand and the deployment of a complementary technology (light-water reactors), which were uncertain when the breeder reaction program began.

The strongest technical and economic argument in favor of breeder reactors stemmed from their ability to generate more fissile material than they used due to their high neutron economy (i.e., the ratio of new fission isotopes per reaction exceeds one). Analysts at the program's inception viewed the reactor as a long-run solution to perceived constraints on uranium supply. Ultimately, the commercial prospects and

[12] An in-depth analysis of the breeder reactor program can be found in Chapter 9 of Cohen and Noll (1991).

economic attractiveness of breeder reactors were strongest when the uranium price escalated enough to offset higher capital costs relative to light-water reactors.

Ultimately, three linked problems eroded the economic rationale for the breeder reactor program. First, uranium discoveries and new approaches for enrichment drove down fuel costs during the lifetime of the project. Meanwhile, construction cost escalation meant that these fuel costs were a smaller fraction of the reactor's lifetime costs. These factors alone dimmed the prospects for breeder deployment. Second, demand for conventional light-water reactors dropped precipitously despite lower fuel costs.[13] This downturn decreased the commercial desirability of liquid metal fast breeder reactors and the value of their associated R&D program. A final problem for the breeder reactor was the decline in demand during the early decades of the program, which delayed new construction of all generators and specifically reduced new orders for nuclear capacity.

Benefit-cost analyses of the project considered a limited range of values in their sensitivity analyses but mainly relied on a deterministic capacity expansion model to assess the breeder's commercial desirability with point-estimates "considered to be most likely" at the time (Cohen and Noll, 1991). Ultimately, the economic attractiveness of breeder reactor R&D was compromised by external conditions and not by the R&D efforts themselves. In retrospect, R&D investment decisions may have been avoided if decision-makers were better equipped with portfolio analysis tools that more carefully provided decision support under multiple simultaneous uncertainties.[14]

[13] The uranium price depends jointly on uranium ore reserves and demand, which are influenced by the deployment of light-water reactors. The market penetration of these reactors depend on the relative economics of nuclear and other competing generation options as well as by electricity demand projections.

[14] This conjecture is conditional on decisions being made on the best-available advice from modeling efforts. However, the analysis in Cohen and Noll (1991) suggests that breeder reactor decisions were clouded by political incentives and that the program's cancellation date extended beyond recommendations from benefit-cost studies with the most up-to-date information.

6.2 Analytical Framework

6.2.1 Assumptions

This work adopts a high-level characterization of optimal aggregate investments in R&D programs.[15] Strategies make allocative decisions for programs with longer-term horizons using a social perspective rather than an individual firm's perspective.[16] The perspective of the public R&D manager (i.e., technology strategist) does not necessarily mean that R&D activities are undertaken by the government. Instead, this frame suggests that R&D efforts produce knowledge that can be considered a public good (i.e., non-rivalrous and non-excludable). The model implicitly assumes that there are no proprietary research boundaries for applying research outcomes, as the publicly disseminated knowledge generated through innovation efforts is an undifferentiated commodity that can be applied by all firms.

As discussed in Section 6.1.3, the operationalization of R&D success in this framework means that investments in R&D programs do not guarantee specific technological outcomes.[17] Additionally, this probabilistic framework differentiates between the benefits of advanced technologies and the benefits of the R&D that acts on these energy technologies.[18]

For the R&D investment uncertainty, the outcome of program efforts is an endogenous function of investment. R&D influences the probability of success for individual

[15] Allocation decisions are analogous to choices about how to distribute wealth across a portfolio of financial instruments.

[16] The hypothetical optimizer has a narrow definition of social welfare. In this context, welfare impacts only account for costs that accrue directly to firms in the electric power sector. This characterization likely biases the benefits of R&D downward.

[17] Other R&D analyses often assume that technological progress has "a strictly positive welfare impact" (Blanford, 2005). Here, R&D success always has a positive cost impact in expectation but not necessarily in every state of the world. For instance, if a successful R&D program shifts the distribution of nuclear capital costs to the left, this optimism about the future of nuclear may lead to hedging behavior that increases anticipated investments in these units. However, high capital costs may obtain, which could leave the decision-maker worse off than in the scenario without these investments. Thus, R&D fundamentally changes expectations about future technological states and typically improves welfare but not in every possible state.

[18] This relationship suggests that the benefits of R&D success do not continue in perpetuity, which is often neglected in the evaluation of actual research programs (National Research Council, 2007; Cohen and Noll, 1991).

programs. The R&D strategist can influence the likelihood of success by adjusting investment levels. In this framework, uncertain returns to investment are captured through the innovation production function (Blanford, 2009). This function is a probabilistic mapping from R&D investment decisions to technological outcomes and is everywhere nonnegative. The key characteristic of the innovation production function is that it exhibits decreasing returns to scale, which means the relationship between R&D investment and probability of success is concave. This functional form does not necessarily imply decreasing returns between investment and value of success, since the value of technological success may be nonlinear and may exhibit increasing returns. The shape of the innovation production function depends on a host of program-specific factors like existing knowledge about the underlying phenomenon, extent of previous research, and potential for technological improvement.

Although the model provides an internally consistent method for translating technological developments into monetary benefits, the valuation of outcomes only captures the impact of R&D on the US electric power sector. The true value of technological developments is likely influenced by many economic factors not accounted for in the model:[19]

- More extensive microeconomic and macroeconomic effects (e.g., consumer surplus gains from lower energy prices; spillovers and benefits to other areas of the energy sector)
- Induced technical change (e.g., learning-by-doing effects)
- International spillovers
- Risk aversion
- Possibility for basic research breakthroughs
- Additional inertia caused by capital stock turnover

[19] The noneconomic impacts of R&D are not captured in this simplified model. Spillover effects are difficult to conceptualize, let alone to quantify. Additionally, the model ignores the many other complex incentives that exist in this heavily regulated industry, which may be mollified or exacerbated by R&D.

- General equilibrium effects (e.g., crowding out of other R&D activities or distortionary effects of taxes)[20]

- Cost-benefit setting

Although the effect of incorporating some of these characteristics into the model is ambiguous, the omission of these effects on the whole likely biases the model's allocation levels downward.

The primary model output is the composition of the R&D investment portfolio. Although the model results suggest the optimal extent of investment, the limitations above indicate that greater emphasis should be placed on the distribution of investment under different budget constraints and conditions of uncertainty. Although recommendations for the total level of R&D allocation are likely underestimates (perhaps severely so), the benefits of some of the successful research programs suggested here are larger than current public expenditures on R&D. Instead of being viewed as the optimal level of investment, the total magnitude of R&D expenditures from this model can be interpreted as a credible lower bound on the value of R&D.[21]

6.2.2 Model Formulation

These assumptions are embedded in the R&D portfolio optimization model, which determines the optimal level of R&D investment in individual technological programs subject to a possible budget constraint. Notation for the mathematical representation of this problems is:

[20] Although the model's partial equilibrium framework cannot capture crowding out effects, such costs could be incorporated into the innovation production functions.

[21] Much like forecasts in the 1986 NASA/WMO assessments for ozone depletion (Parson, 2003), these values may constrain the bargaining range for policy proposals.

Notation

n	number of R&D programs
α	R&D investment allocation vector; $\alpha = \{\alpha_1, \ldots, \alpha_n\}$
$\widetilde{\theta}$	random variable for the R&D program state
$\theta \in \Theta$	R&D program outcome
θ_0	baseline technological state with no R&D success
$\widetilde{\omega}$	random variable for market conditions
$\omega \in \Omega$	market outcome
$\mathbf{p}(\alpha)$	probability density function with multiple technologies
p_i	innovation production function for program i; $p' > 0$, $p'' < 0$
$S(\theta)$	set of technologies with R&D successes for outcome θ
B	budget
$V(\theta, \omega)$	present value of R&D success

Investments can be allocated across a range of technological R&D programs, and decision variables in this allocation vector α are continuous in the model. Definitions for the individual R&D programs included in the model are found in Section 6.2.3.

The function $V(\theta, \omega)$ determines the value of R&D success using a stochastic diffusion model. This stochastic optimization framework utilizes the two-stage capacity planning and dispatch model from Chapter 3 and adopts the cost-minimizing objective of utilities and generators. The model is unique in its ability to represent the uncertain process of diffusion explicitly. This capability is important in decision contexts with irreducible elements of risk, which require decisions to be made before uncertainty is resolved.[22]

The R&D program outcome θ can be thought of as a set of parameters for the joint distribution over the second-stage technological state. The state of the world ω represents the realization of all uncertainties that are exogenous to the technological state (e.g., climate policy stringency). Each R&D program has two possible outcomes, which gives rise to a technological baseline distribution when the program fails and

[22]In order to reduce the dimensionality of the optimization problem, the only exogenous market uncertainties considered in this chapter are the carbon tax stringency, natural gas prices, and public acceptance of CO_2 storage, which were found to be the most significant uncertainties in Chapter 4.

to an enhanced R&D distribution when the program is successful. Thus, the outcome space for $\widetilde{\theta}$ has a total of 2^n elements. $S(\theta)$ can be conceptualized as the production-possibility set associated with R&D program outcome θ so that, in the advanced technological state θ', $S(\theta_0) \subset S(\theta')$.

General Optimization Problem

The R&D portfolio manager's objective is to maximize the expected net benefit (i.e., gross benefits less expenditures) of investments subject to a budget constraint. This net benefit is equal to the expected discounted value under a range of economic, policy, and technological scenarios. Assuming risk neutrality, the technology strategist's optimization problem is:

$$\max_{\alpha} \mathbb{E}\left[V(\widetilde{\theta}, \widetilde{\omega})\right] - \sum_{i=1}^{n} \alpha_i \tag{6.1}$$

$$\text{s.t.} \sum_{i=1}^{n} \alpha_i \leq B, \;\; \alpha_i \geq 0 \tag{6.2}$$

Incorporating the probability density function $\mathbf{p}(\alpha; \theta)$ explicitly (for each R&D program outcome θ with allocation vector α) expands the first term in Equation 6.1:

$$\mathbb{E}\left[V(\widetilde{\theta}, \widetilde{\omega})\right] = \sum_{\theta \in \Theta} \mathbf{p}(\alpha; \theta) \mathbb{E}\left[V(\theta, \widetilde{\omega})\right] \tag{6.3}$$

If λ_B is the shadow price on the budget constraint, the first-order conditions for this problem become:

$$\sum_{\theta \in \Theta} \frac{\partial \mathbf{p}}{\partial \alpha}(\alpha; \theta) \mathbb{E}\left[V(\theta, \widetilde{\omega})\right] = 1 + \lambda_B \tag{6.4}$$

$$\lambda_B \left(\sum_{i=1}^{n} \alpha_i - B \right) = 0 \tag{6.5}$$

where the first equation represents stationarity and the second represents complementary slackness. The left-hand side of Equation 6.4 can be viewed as the benefit of a marginal dollar of R&D investment (i.e., the marginal increase in the expected future welfare via an incremental change in investment). The right-hand side of Equation 6.4 expresses the opportunity cost of investment.

Quantifying the Expected Value of Success

There are a few approaches to quantifying the expected value of R&D program success $\mathbb{E}[V(\theta, \widetilde{\omega})]$ given an outcome θ. Assuming the objective function $f(\cdot)$ maps decisions to costs, the wait-and-see approach (as used in Blanford (2009); Anadon et al. (2011); and other uncertainty propagation frameworks) quantifies the success valuation as:

$$\mathbb{E}\left[V_{ws}(\theta, \widetilde{\omega})\right] = \mathbb{E}_{\omega}\left[\min_{x} f(x; \theta_0, \omega)\right] - \mathbb{E}_{\omega}\left[\min_{x} f(x; \theta, \omega)\right] \tag{6.6}$$

The here-and-now approach (as used here) quantifies the success valuation as:

$$\mathbb{E}\left[V^{*}(\theta, \widetilde{\omega})\right] = \min_{x} \mathbb{E}_{\omega} f(x; \theta_0, \omega) - \min_{x} \mathbb{E}_{\omega} f(x; \theta, \omega) \tag{6.7}$$

The stochastic value function $V^{*}(\theta, \widetilde{\omega})$ comes from the two-stage stochastic programming model described in Chapter 3. The model is characterized by its sequential decision-making structure and dependence on energy system inertia.[23]

Single Technology Case

Given an outcome space with two elements (i.e., for the baseline technological state and R&D success state), $p(\alpha)$ is the probability of the single program's success. The first-order condition in Equation 6.4 becomes:

$$\frac{\partial p}{\partial \alpha} = \frac{1 + \lambda_B}{\mathbb{E}[V(\theta', \widetilde{\omega})]} \tag{6.8}$$

where θ' represents the advanced technological (success) state for the R&D program.

[23]This trait is a key benefit of using a detailed, bottom-up capacity expansion model. It represents existing assets, capital vintaging, expansion constraints based on existing pipeline projects and resources, as well as a range of available technologies.

These conditions offer insights about the R&D investment problem. As described in Blanford (2006), the optimal R&D investment level occurs when the marginal probability $\frac{\partial p}{\partial \alpha} \equiv p_\alpha$ equals the threshold quantity on the right-hand side of Equation 6.8 when p_α is strictly decreasing. In the case where the marginal probability is less than the threshold at $\alpha = 0$, the program should not be funded. In the case where the marginal probability exceeds the threshold at $\alpha = B$ with $\lambda_B = 0$, the budget constraint is binding, and the optimal solution is to invest the entire budget in the program.[24] Otherwise, an interior solution for α^* exists.

The value of the left-hand side depends on characteristics of the innovation production function. The shape of this function (and consequently of the marginal probability curve) depends on characteristics of the R&D program and how amenable the technology is to development. In contrast, the right-hand side threshold value depends primarily on deployment-related factors of the technology and of the markets into which it will diffuse. Ultimately, the optimal investment in a specific R&D program balances the expected market value of the program's success with the program's likelihood of success conditional on R&D expenditures.

Portfolio Analysis

If programs are independent, the portfolio optimization problem reduces to the individual technology case, and the problems can be solved in isolation to develop unilateral strategies for each R&D program. However, most R&D decision contexts cannot be modeled under this assumption due to:

1. *Market Competition*: Interactions between complementary and substitute technologies in the energy marketplace influence the value of R&D success of specific technological programs.

2. *Spillovers*: Innovation production functions may be correlated, which influences the probability of success for R&D programs.

[24]The shadow price should rise such that $p_\alpha(B)$ equals the right-hand side threshold. In cases where the budget constraint is binding, λ_B decreases in B.

This work focuses on program interactions to capture diffusion-related uncertainties and not on spillovers, which is an important area for future research. For the case without spillovers, the probability density function with multiple technologies is:[25]

$$\mathbf{f}(\alpha;\theta) = \prod_{i=1}^{n} \left[\mathbb{1}_{\{i\in S(\theta)\}} p_i(\alpha_i) + \mathbb{1}_{\{i\notin S(\theta)\}}(1-p_i(\alpha_i))\right] \tag{6.9}$$

where the left-hand side represents the probability of technological outcome θ, and the right-hand side terms represent the possibility that program i succeeds in outcome θ and fails, respectively.

The optimality conditions can be expressed as:

$$\sum_{\theta\in\Theta} \frac{\partial p_i}{\partial \alpha_i}(\alpha_i;\theta)\mathbb{E}\left[V(\theta,\widetilde{\omega})\right] = 1 + \lambda_B \quad \forall i \tag{6.10}$$

$$\lambda_B \left(\sum_{i=1}^{n} \alpha_i - B\right) = 0 \tag{6.11}$$

These conditions comprise a system of simultaneous equations. Since the valuation of R&D success is dependent on the outcomes of complementary and substitute technologies and consequently investment across other R&D programs, this system of equations must be solved simultaneously to develop the optimal portfolio. This system equates the success probability of the marginal investment across R&D programs for all technologies, which means that the portfolio optimality conditions equate marginal productivity across all program investments.

6.2.3 Specification and Data

The four energy technology R&D programs incorporated in this research reflect the focus on the electric power sector. This concentration is motivated by the industry's status as the largest greenhouse gas emitter in the US economy and its low R&D intensity. Although these four programs do not cover the full portfolio of potential R&D projects, they are broadly representative of different mechanisms for reducing

[25] The innovation production function decomposition for program i is expressed as $p_i(\alpha_i)$.

the emissions intensity of production in the electric power sector.[26]

1. **Natural gas turbine efficiency (GAS):** Improvements in the first-law efficiencies of fossil-based units, particularly for natural gas, can take many forms, including improvements to existing configurations (e.g., more efficient combustion or the introduction of high-temperature materials) or the diversification into new architectures (e.g., fuel cell and gas turbine hybrid systems). No matter the manifestation, such advances increase the production from a given exergetic input. The associated reduction in emissions per unit output makes it more attractive to construct and/or operate these units even under moderate climate policy conditions.

2. **Carbon capture and sequestration (CCS):** The CCS family of technologies interacts with other fossil-based assets.[27] The CCS R&D program considered here includes both coal and natural gas with capture systems and assumes perfect spillovers between programs. Unlike its fossil-based counterparts with capture, CCS technologies likely cannot be profitability deployed in a policy environment without a carbon price. However, unlike renewables, the profitability of CCS also is threatened under an extremely stringent climate policy regime, as discussed in Chapter 4. Thus, the diffusion of CCS and the impacts of its R&D outcomes are sensitive to the climate policy uncertainty.

3. **Nuclear (NUC):** Like renewables, nuclear power is a low-carbon technology. For this reason, nuclear R&D programs that reduce capital costs are most valuable under stringent climate policy conditions. However, due to its low fuel cost volatility relative to fossil-based units, nuclear can be economically competitive even without a price on carbon, as suggested by its prominent role as a hedging technology in Chapter 4.

[26] The model structure and selected programs implicitly assume that technical and cost advances improve known technologies. Such incremental technological change has historically driven innovative activity in this sector (Alic, Mowery, and Rubin, 2003).

[27] Generally, CCS technological developments may reduce capital costs, decrease parasitic energy losses from capture equipment, or provide cheaper transportation and geologic storage. Here, R&D programs focus only on capital cost reductions.

4. **Solar (SOL):** Although many renewable technologies may gain larger market shares in the coming decades (e.g., solar PV and thermal, wind, biomass, fuel cells, geothermal), the R&D work here includes only utility-scale solar. The R&D program is intended to reduce the capital cost of the technology per unit output. A success in this program may potentially lead to increased deployment even in the absence of a climate policy; however, solar's greatest comparative advantage occurs when climate policy targets are most stringent.

Due to the structure of the expert elicitations underlying values in the model, R&D activities influence cost parameters of the CCS, nuclear, and solar technologies but affect the performance of gas-turbine-based technologies.

As described in Chapter 3.4, technological outcomes are characterized through discrete distributions based on expert elicitations for individual energy technologies. Uncertainties for coal with CCS, gas with CCS, nuclear, and solar are incorporated in the model as distributions over investment costs. These distributions come from expert elicitations conducted at the Harvard Kennedy School (Anadon et al., 2011). The natural gas combined cycle efficiency uncertainty comes from the elicitations described in Chapter 3.4 and Appendix B.

Like Blanford (2009), the specific form for the innovation production function of R&D program i is a bounded exponential:

$$p_i(\alpha_i; \rho_i, \beta_i) = \rho_i \left(1 - e^{-\frac{\alpha_i}{\beta_i}}\right) \tag{6.12}$$

where $0 < \rho_i < 1$ and $\beta_i > 0$. This form satisfies the decreasing returns to scale assumption ($\frac{\partial^2 p_i}{\partial \alpha_i^2} < 0$) so that the marginal probability is strictly decreasing in own-program investments. Additionally, $p_i(0) = 0$ and $\lim_{\alpha_i \to \infty} p_i(\alpha_i) = \rho_i < 1$, which makes the bounded exponential form feasible for a probability mapping. The parameter ρ_i represents the asymptotic limit of the success probability for a program as R&D investment increases. The parameter β_i scales the innovation production function and determines the rate at which the probability of success approaches its limiting value ρ_i.

Although this functional form is based on Blanford (2006), the parameterizations for ρ_i and β_i used here come from actual expert elicitations instead of *ad-hoc* estimates. These program-specific values are outputs of elicitation-based analyses on CCS (Baker, Chon, and Keisler, 2009b), nuclear (Baker, Chon, and Keisler, 2008), and solar (Baker, Chon, and Keisler, 2009a).[28] Figure 6.5 shows the resulting innovation production functions from the nonlinear least squares analysis.

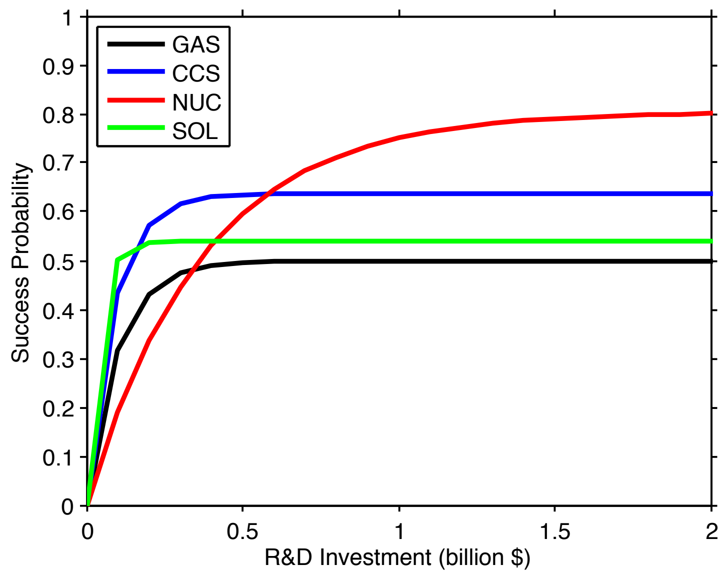

Figure 6.5: Innovation production functions for the natural gas (GAS), carbon capture (CCS), nuclear (NUC), and solar (SOL) R&D programs.

6.3 Results

6.3.1 Valuation of Program Outcomes

This section explores the expected value of R&D program successes before determining optimal portfolio investments. The two-stage stochastic programming model was

[28]Outputs from these papers show how the probability of success changes as a function of R&D investment. It is important to note that the conceptualization of success used in the Baker papers differs from the one used here. This difference suggests that future elicitation and analysis work should consider how to improve parameterizations of innovation production functions under alternate conceptions of R&D success.

run for all $2^n = 16$ possible combinations of program successes and failures under uncertainties in the carbon tax stringency, natural gas prices, and public acceptance of carbon dioxide (CO_2) storage. As described in Section 6.2.2, the expected value of a specific technological outcome is defined as the difference between the expected cost in that particular scenario and the baseline technological scenario.

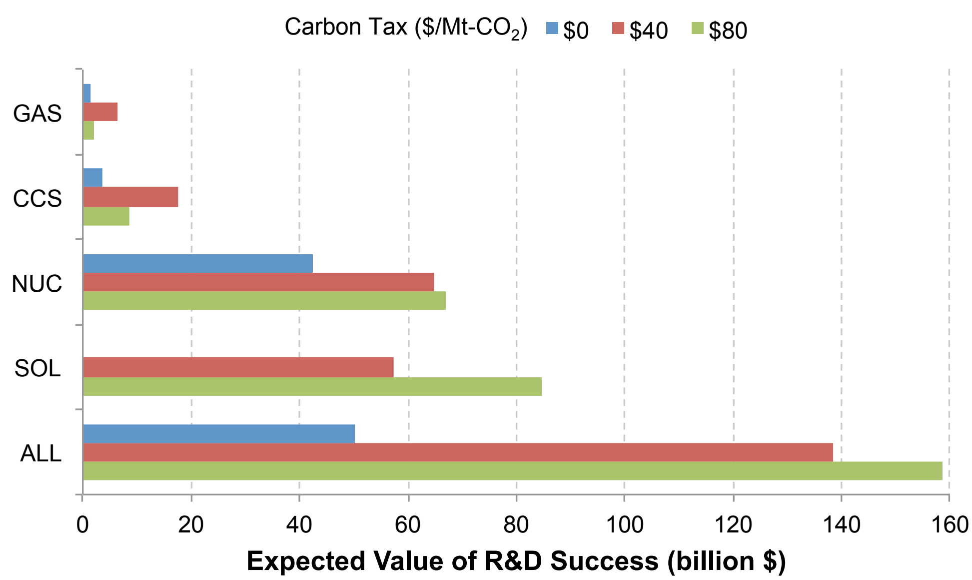

Figure 6.6: Expected value of R&D program success (billion $) decomposed by the outcome of the climate policy uncertainty.

Figure 6.6 shows the value of success conditional on the outcome of the climate policy uncertainty.[29] The top of the figure shows values when the R&D programs are considered individually, and the bottom cluster represents the joint value when all programs succeed. In general, the expected value of R&D is extremely sensitive to the stringency of climate policy, which makes explicit consideration of this uncertainty important given how its future state cannot be known when R&D decisions are made.

[29]These results are conditioned both on the stochastic model and on the description of research success from the elitications. As pointed out by Blanford (2009), these values would be considerably higher if the alternative to program success were that the technology is never available. Although this effect can be captured in this framework, the elicitations do not indicate that it is likely for any of the technologies considered here.

For the gas and CCS programs, R&D successes are most valuable under moderate carbon taxes due to the fact that they permit greater generation and deployment from lower-cost fossil units. However, the \$80/Mt-CO_2e carbon tax is sufficiently stringent to deter investments in fossil units under many scenarios. For the low-carbon programs (i.e., nuclear and solar), the R&D value increases in the realized carbon tax. When these carbon tax scenarios are weighted by their respective probabilities, the nuclear R&D program has the highest value of success due to its optionality and lower carbon intensity. Nuclear's high expected value of R&D success reflects its broader base for diffusion under a range of market conditions compared with other technologies, which are only competitive in more limited domains.

Another conclusion is that the expected value of R&D success is slightly subadditive (i.e., the value of the joint scenario is less than the sum of the individual successes). On one hand, the complementary interaction between the gas and CCS programs makes their value of R&D success superadditive. On the other hand, substitution effects between low-carbon technologies in the marketplace slightly outweigh this effect and lead to a net subadditive value of success for all of the programs.

6.3.2 Optimal R&D Portfolio Investments

Using the R&D decision model and innovation production functions described in Section 6.2, the valuations of R&D outcomes in the previous section can be used to construct optimal portfolio investments. Figure 6.7 compares R&D program allocations when decisions are made for programs individually and when the optimal portfolio is chosen.[30] Analyzing individually optimal R&D investments (Equation 6.8) can identify the influence of the innovation production function on allocation decisions without considering program interactions or spillovers. On the other hand, optimal portfolio allocation levels incorporate competition between R&D technological outcomes in the marketplace (Equation 6.10).

[30] R&D investments in Figure 6.7 are shown on an annual basis for the assumed ten-year durations of the programs.

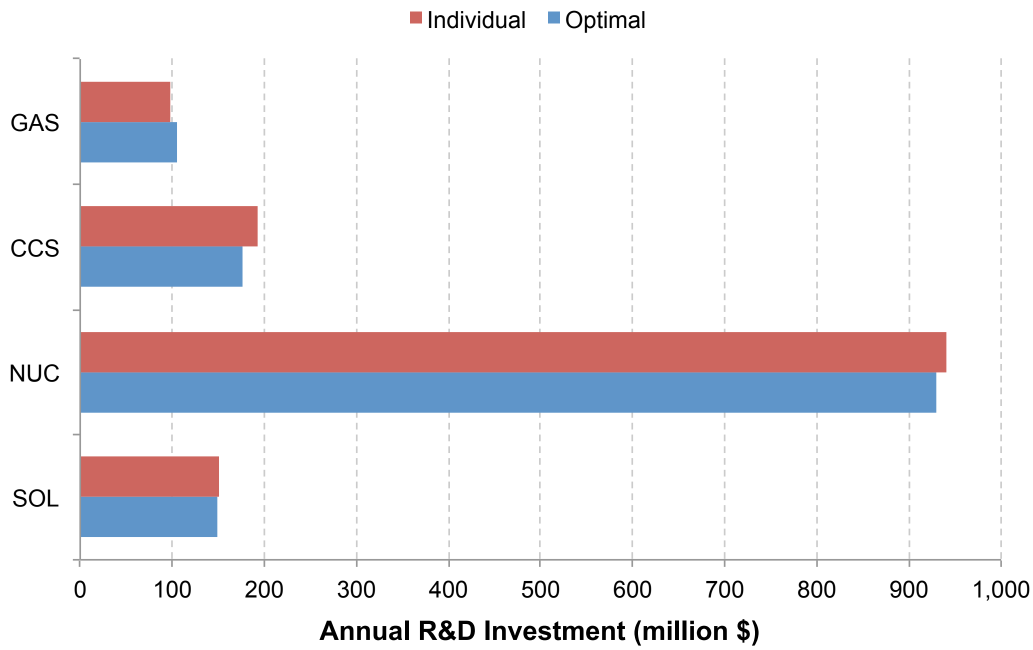

Figure 6.7: Comparison of R&D allocation (million \$) across programs under individual and optimal portfolio compositions.

Figure 6.7 suggests that optimal allocation levels are similar when R&D programs are considered individually or jointly. Portfolio investments in programs are typically lower due to market competition when programs are successful. Investments in the gas efficiency program increase in the joint portfolio case due to the complementarity between gas and CCS R&D successes, as efficiency improvements in the base natural gas combined cycle design improve the efficiency of CCS-equipped units as well.

The highest recommended investments are for the nuclear R&D program, as allocations for nuclear exceed all other programs combined. Despite its low success probability for small investments, the nuclear program has highest limiting probability for large investments and the largest value of success across a range of scenarios. Although solar has high value of success and high marginal returns for early investments (i.e., a β parameter that is consistent with a lower degree of initial difficulty), it also has a smaller limiting probability, which discourages greater investment. This result underscores the importance of the ρ parameter in the innovation production function, which can be interpreted as a proxy for the level of confidence that large R&D investments will result in a favorable outcome.

In many contexts, budget constraints are binding for R&D portfolio allocation decisions. Figure 6.8 demonstrates how the optimal portfolio composition changes as the annual R&D budget decreases. As the budget tightens, the portfolio composition gradually becomes less diverse, with nuclear and solar investments crowding out gas and CCS. The high success valuations for the nuclear and solar programs are complemented by their favorable R&D program characteristics, as the nuclear program has a high limiting success probability (due to ρ in the innovation production function) and solar has high marginal returns for initial investments (due to β).

The homogeneity of the portfolio composition across a range of assumptions highlights the importance of the structure of the R&D valuation model. This conclusion is reinforced by the large share of nuclear, which occurs in spite of the comparatively unfavorable returns for initial program investments.

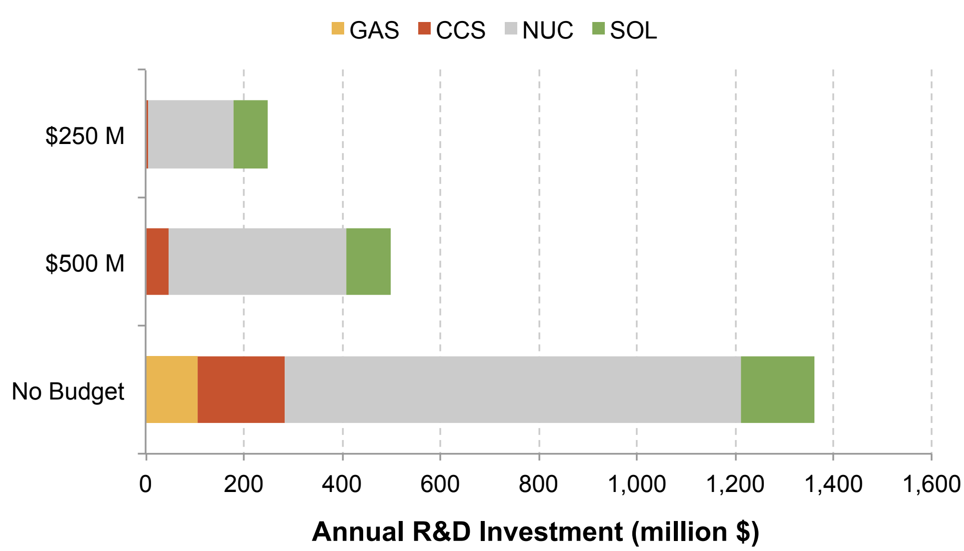

Figure 6.8: Optimal portfolio composition under different budgets.

It can be shown through comparative statics that stronger complementarity between programs leads to more evenly distributed allocations, whereas stronger substitutability leads to a larger share of the dominant technology (Blanford, 2006).

However, it is important to note that the portfolio composition is never completely one-sided, even under the restrictive budget constraints. Portfolio diversification in this setting is not motivated by risk aversion, as in portfolios of financial securities (Markowitz, 1952), but by a combination of other factors in the risk-neutral technology strategist's optimization problem (Blanford, 2009):

- **Decreasing returns to scale:** The marginal productivity of an R&D program decreases in investment. Due to the optimality condition of equal marginal returns across investment alternatives, R&D expenditures are likely to be moved toward programs that offer the highest marginal returns. Thus, it is unlikely that a single program will dominate the portfolio, as investments are spread to exploit the most productive range for each program.

- **Uncertainty:** Diversification can reduce risk under conditions of simultaneous uncertainty in program outcomes and exogenous market conditions. This risk management function of diversification provides insurance against the possibility that an individual program does not resolve as expected.

- **Heterogenous diffusion markets:** Portfolio diversification may occur in situations with many applications and diverse markets for energy technologies associated with R&D investments.

6.3.3 Valuation in Alternate Decision-Making Contexts

As discussed in Section 6.1.4, standard models value R&D using the wait-and-see (i.e., perfect information) approach, which assumes that uncertainty about market conditions is resolved at the beginning of the time horizon. One contribution of this work is to incorporate uncertainty about these conditions explicitly in the planning process. A third approach is to replace the stochastic parameters by their expected values and to plan according to the resulting deterministic problem (i.e., the expected-value approach). This section compares the expected value of R&D success under these three alternate decision-making contexts.

As shown in Appendix C, a related theoretical result is that the effect of the decision-making approach on R&D valuation is equivocal. One implication is that modeling efforts using a sequential decision-making framework are important to reduce the ambiguity of this effect in a particular setting. Since wait-and-see approaches are typically used for R&D valuation in the literature (Anadon et al., 2011; Blanford, 2009), the conclusion here underscores the importance of understanding the limitations of results from models with simplified treatments of uncertainty. In other words, such models do not necessarily provide upper bounds on R&D values, since the bias introduced by adopting a wait-and-see approach can be upward or downward depending on characteristics of the model.

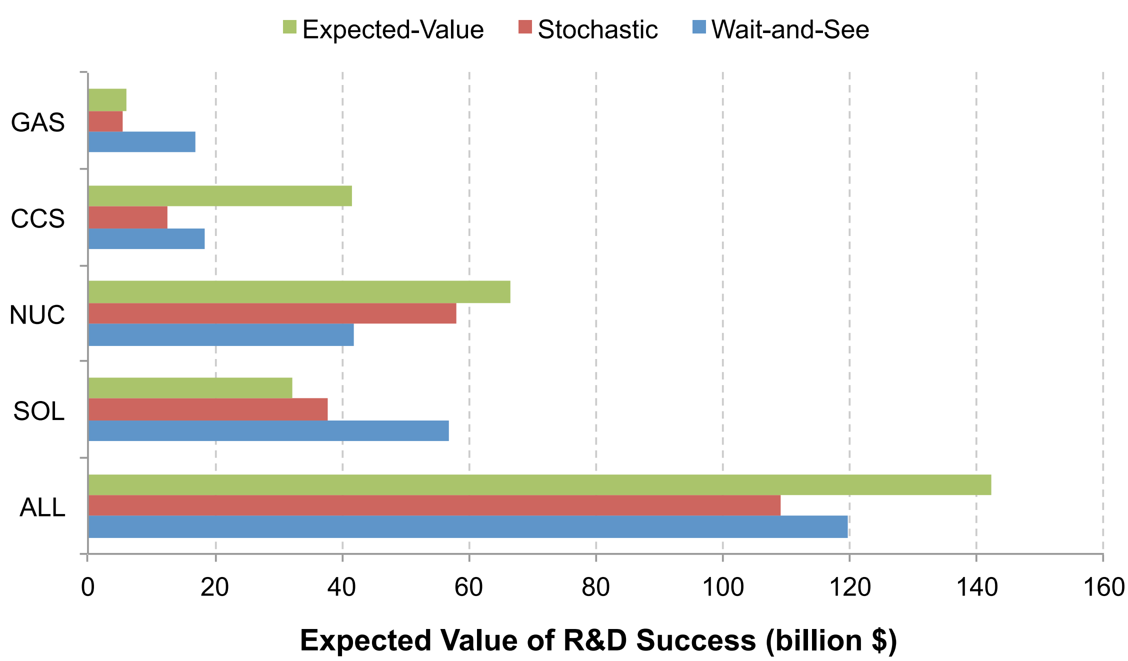

Figure 6.9: Comparison of expected R&D success valuations (billion \$) under alternate decision-making approaches.

For the numerical simulations here, the most surprising result is that the R&D success value is largest under the expected-value approach in three of the five cases, as shown in Figure 6.9. Recall that the Chapter 4 results suggest that, relative to the other decision contexts, planning under the expected-value approach neglects many contingencies and gives rise to greater stranded investments once uncertainties are resolved. When R&D programs are successful, the expected-value strategy mobilizes and deploys technologies at lower costs in the second stage. The reason for the greater R&D value in the expected-value context is that suboptimal first-stage decisions lead to the eventual decommissioning of some of these units under certain scenarios once information is available about fuel prices and the policy environment. When substantial construction occurs to replace these assets, R&D has lowered costs for advanced technologies. Thus, the benefits of R&D are larger under these scenarios for suboptimal decision-making approaches compared with here-and-now and wait-and-see approaches, where robust planning precludes the need for substantial corrective actions in the second stage.[31] This conclusion is particularly relevant given that utility resource planning currently relies on relatively simple approaches for uncertainty analysis and does not model uncertainty explicitly (i.e., using sequential decision-making frameworks), as described in Chapter 3.4.

The primary takeaway is that R&D can be more valuable in second-best planning environments. Previous research indicates that R&D is more valuable in second-best climate policy environments (Baker and Solak, 2013), but the research here indicates that R&D may also exhibit greater value in another type of suboptimal setting—namely, contexts in which capacity planning decisions are made using an expected-value approach to decision-making under uncertainty. R&D investments act as insurance against suboptimal decisions, assuaging their negative cost impacts.

Comparing the stochastic and wait-and-see approaches in Figure 6.9 suggests that the perfect information approach yields a higher expected value of R&D success for many of the technologies. In these cases, the wait-and-see approach makes greater capacity investments in the second stage than the stochastic approach, which can take advantage of the reduced capital costs made possible by R&D successes. In contrast, the value of a nuclear R&D success is larger under the stochastic planning approach

[31] One caveat with this conclusion is that the runs assume that second-stage decisions adjust to updated information and make optimal investments accordingly. However, the same conditions that gave rise to suboptimal decision-making in the first stage may hinder the firm's ability to properly consider the best-available information and to make informed decisions in later stages as well.

due to its prominent role as a hedging technology.[32] The larger expected value of R&D success for the stochastic approach occurs primarily for low realizations of the carbon price uncertainty. This expanded rollout of hedging capacity would not have otherwise been built if perfect information were available about carbon policy. Overall, this comparison indicates that traditional, wait-and-see R&D valuation approaches likely overestimate the expected value of R&D success but undervalue the optionality and hedging potential of technologies like nuclear relative to the here-and-now approach.

6.3.4 Alternate Investment Environments

The results in earlier sections assume that exogenous market uncertainties like future climate policies and natural gas prices are not known when R&D portfolio decisions are made. This section investigates how R&D portfolio allocations would change in various investment environments if the second-stage conditions were known *ex ante*. Figure 6.10 shows heatmaps of optimal annual R&D investments for all four programs decomposed by the realized climate policy and natural gas price.

For the gas efficiency program, the optimal R&D investment is heavily influenced by natural gas prices. When gas prices are low, substantial R&D is justified due to the lucrative market environment for gas units regardless of the climate policy stringency, since CCS can be employed with gas if the carbon tax is high. When gas prices are high (or modest but with higher carbon taxes), market opportunities for natural gas become scarce, which justifies reallocating R&D funds into other programs if this outcome is known *ex ante*.

The highest investments in CCS occur under moderate carbon taxes. Allocations are asymmetric on either side of this value, as some investment is warranted under the most stringent policy but avoiding investment is optimal for the no-policy scenario. That greater investment occurs when natural gas prices are high is a reflection of increased market opportunities for deploying new coal units with CCS or for retrofitting existing coal facilities.[33]

[32]The expectation of lower capital costs makes nuclear an even more attractive hedging option, as an additional 13 GW of new construction begins before 2025 under the enhanced R&D scenario.

[33]The CCS R&D investment under the no-policy scenario with high gas prices reflects the end

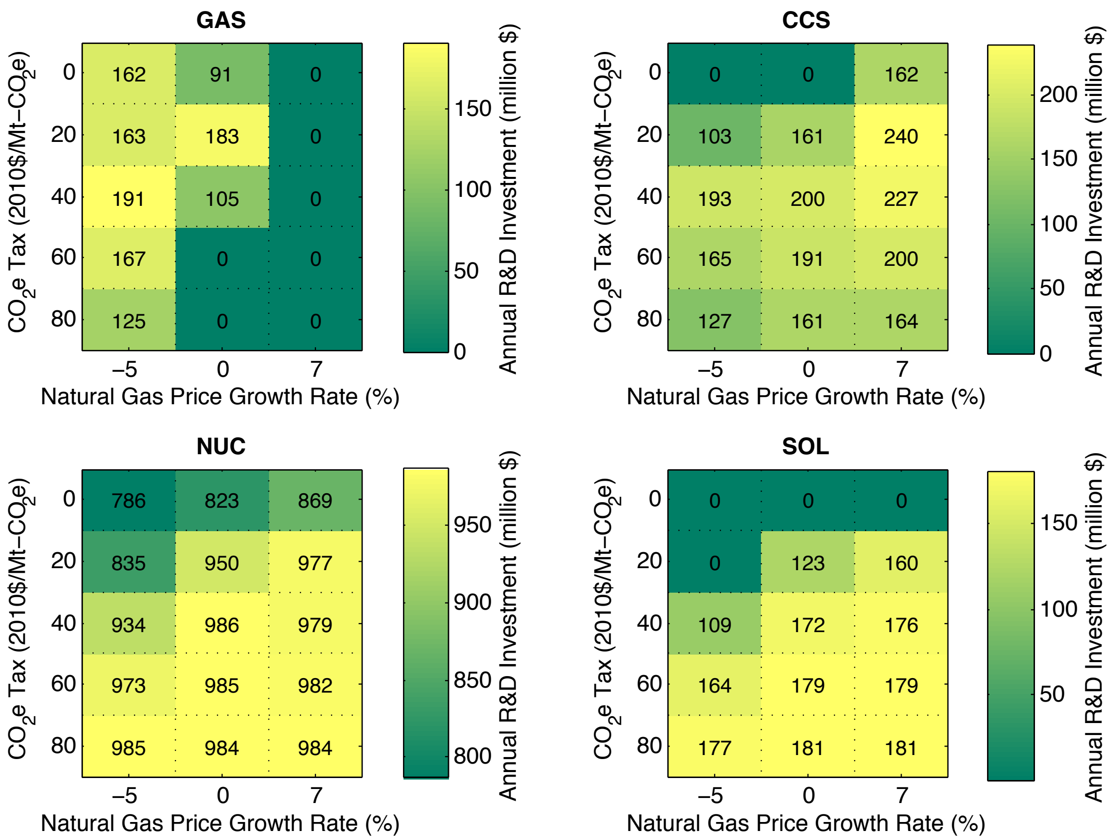

Figure 6.10: Optimal annual R&D investment (million \$) for each technology decomposed by the second-stage climate policy and natural gas price growth rate.

Investments in the nuclear R&D program are considerably higher than for other technological programs under all scenarios. R&D success for this program even has substantial value in the scenario with the least favorable market conditions (i.e., no carbon tax with low natural gas prices), since nuclear capital costs with successful R&D can ensure deployment even without a carbon tax.

Like CCS, market applications for solar are largest with a nonzero carbon tax. The elicitations indicate pessimism about potential cost reductions from utility-scale

effect of retrofitting coal plants with CCS that would otherwise retire during last decade during time horizon. In the model, it is cheaper to retrofit these older units to extend their operating lifetimes rather than to retire and build completely new facilities.

solar in the next few decades.[34] Like nuclear, solar investments are less sensitive to gas prices, except when the carbon tax is low.

The box plots in Figure 6.11 for abatement costs (left) and emissions (right) compares the baseline and R&D scenarios under three different climate policy stringencies. If the policy environment involves a stringent carbon tax, R&D has larger impact on abatement costs and an almost negligible effect on annual emissions. On the other hand, if the policy environment favors lower abatement effort, R&D has a more appreciable impact on emissions and a smaller effect on abatement cost. These trends will be important determinants of illustrating how R&D is more important in second-best policy environments in Section 6.3.5.

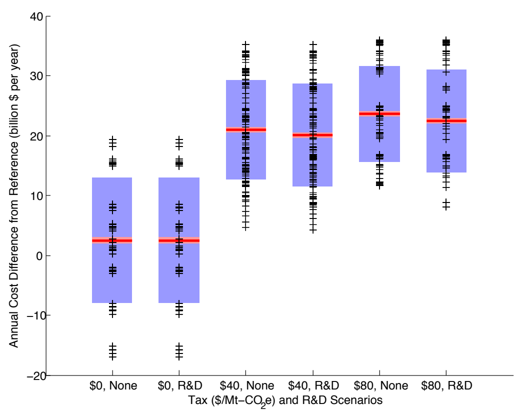

(a) Abatement cost

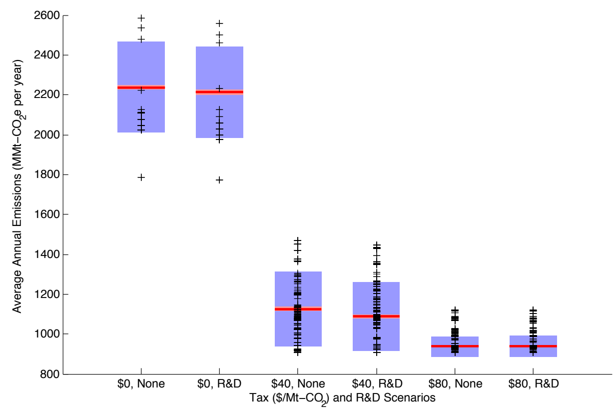

(b) Emissions

Figure 6.11: Box plots of annual abatement cost (left) and emissions (right) comparing the baseline technology case without R&D ("None") and the R&D success case ("R&D"). Each plus sign represents one scenario of the 7,290 possible outcomes. Abatement costs represent annual differences from the reference case without climate policy and all other parameters held at their mean values (e.g., allowing for negative values when gas prices or technological costs are lower than the mean value).

[34]Given the rapid nature of solar developments, these elicitation results must be tempered by an understanding that the cost elicitations are likely outdated and biased upward. Appendix B discusses best practices for expert elicitations and potential biases in existing probabilistic assessments for energy technologies.

One function of R&D is to lower the cost risk associated with climate policy compliance costs. Figure 6.11 suggests that R&D investments provide some degree of risk reduction for cost outcomes, especially for moderate to high carbon taxes. This result recasts the R&D portfolio strategy as a risk management problem.

6.3.5 Suboptimal Climate Policies and R&D Investments

Thus far, the results assume that the benefits of R&D include only direct financial gains for utilities and that the R&D manager makes expenditures that are equal to the recommendations from the model. This section relaxes these assumptions and incorporates the social cost of carbon into the ledger of R&D benefits and then explores the welfare impacts of R&D overinvestment and underinvestment. These analyses answer questions like: What happens when either climate policies or R&D spending are suboptimal? How are program performance and net benefits impacted?

R&D and the Social Cost of Carbon

The earlier results examined the expected value of R&D success without accounting for damages from greenhouse gas emissions.[35] This section explores how these results would change if the so-called social cost of carbon were included in the analysis as additional costs after model runs. The social cost of carbon is uncertain and subject to many assumptions about amplifying feedbacks (e.g., thawing of vast deposits of frozen methane), catastrophic impacts (e.g., slowdown or shutdown of Atlantic Meridional Overturning Circulation), and economic parameters (e.g., social rate of time preference). Therefore, the expected value of R&D is calculated for a range of potential values for the social cost of carbon.

[35] As mentioned in Chapter 3.2, the decision-makers in this formulation are utilities and generators. Their optimization problem minimizes private system costs without accounting for social costs. They only consider climate policy targets to be uncertain, even though this policy may not fully internalize the externalities associated with greenhouse gas emissions.

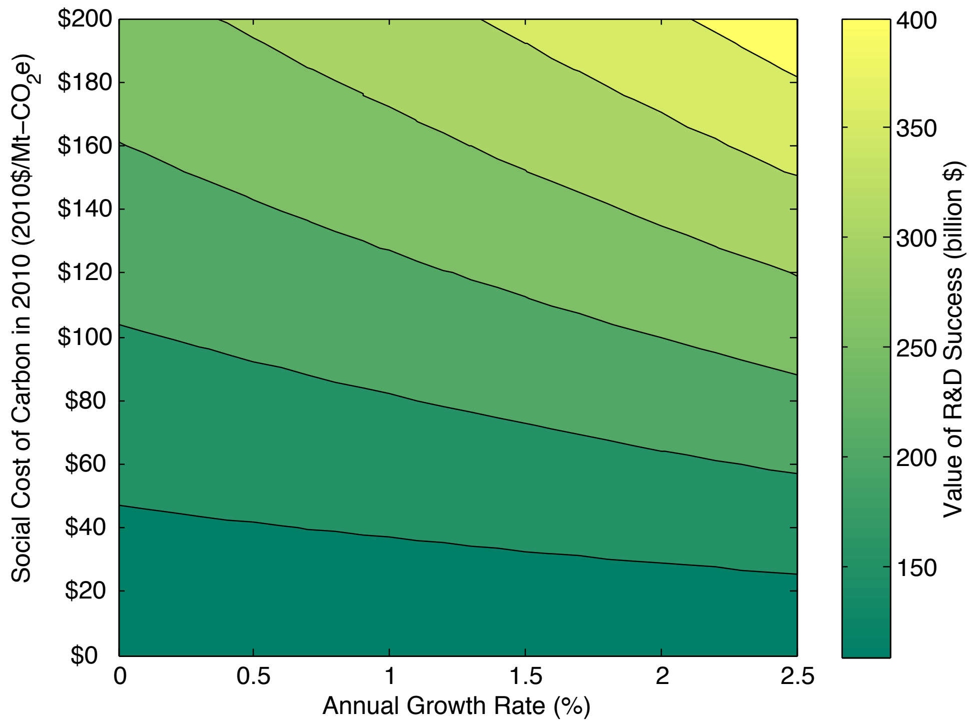

Figure 6.12: Expected value of R&D success (billion $) for all technologies under various social costs of carbon in 2010 and annual growth rates.

Figure 6.12 shows equipotential curves for the expected value of R&D success for all programs at various values for the social cost of carbon in 2010 and annual growth rates over time. The value from Section 6.3.1 of $109 billion (where the social costs of carbon are excluded) is shown at the origin. Including climate damages makes R&D successes more valuable, since R&D helps to mobilize low-carbon technologies that would not otherwise be cost-competitive when policy is less stringent than the socially optimal level.

These results mirror those in Section 6.3.4. Although R&D spending may be valuable as means to correct technological market failures like appropriability, it is *a fortiori* valuable for its ability to ameliorate a negative environmental externality that has not been properly internalized. Baker and Solak (2013) also conclude that R&D is more valuable in second-best policy environments but is referring to second-best policies involving supererogatory abatement. Here, the results suggest that a similar rationale holds for the availability of R&D success under a suboptimal policy milieu, where the carbon tax is less than its socially optimal level. If R&D decisions are made in a suboptimal policy environment with low abatement, R&D has a smaller impact on lowering abatement costs but a larger impact on emissions. Utilities' anticipation of climate legislation (and not its actual realization) drives investments in R&D for low-carbon technologies. When these programs succeed, some of these reduced-carbon-intensity generators will be deployed, even when a stringent climate policy is not implemented. Therefore, the benefits of R&D are higher when the social cost of emissions is incorporated into the accounting.

Incorporating the social cost of carbon into the R&D valuation makes low-carbon technologies especially attractive, as Figure 6.13 illustrates. When the social cost of carbon is \$23/Mt-CO_2e, solar R&D becomes more valuable than nuclear. As the social cost of carbon increases, market opportunities for natural gas and CCS erode, which shifts R&D portfolio allocations toward solar and nuclear.

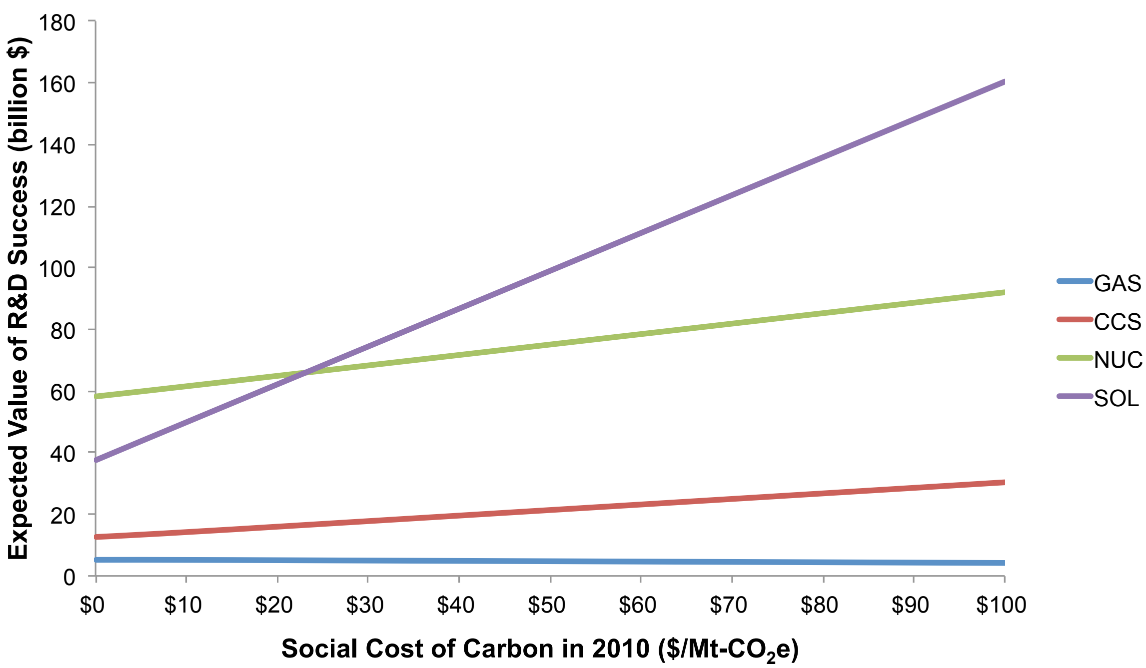

Figure 6.13: Expected value of R&D success (billion \$) by technology under various social costs of carbon in 2010 and annual growth rates.

Suboptimal R&D Budgets

Thus far, the analysis has assumed that the R&D manager makes expenditures that are equal to the recommendations from the model. Here, I investigate the welfare impacts of suboptimal budgets. The optimal unconstrained R&D budget according to Section 6.3.2 is $1.36 billion annually. Figure 6.14 considers three alternate budget levels: no R&D, 2x the optimal R&D budget ($2.72 billion), and 4x the optimal budget ($5.44 billion).

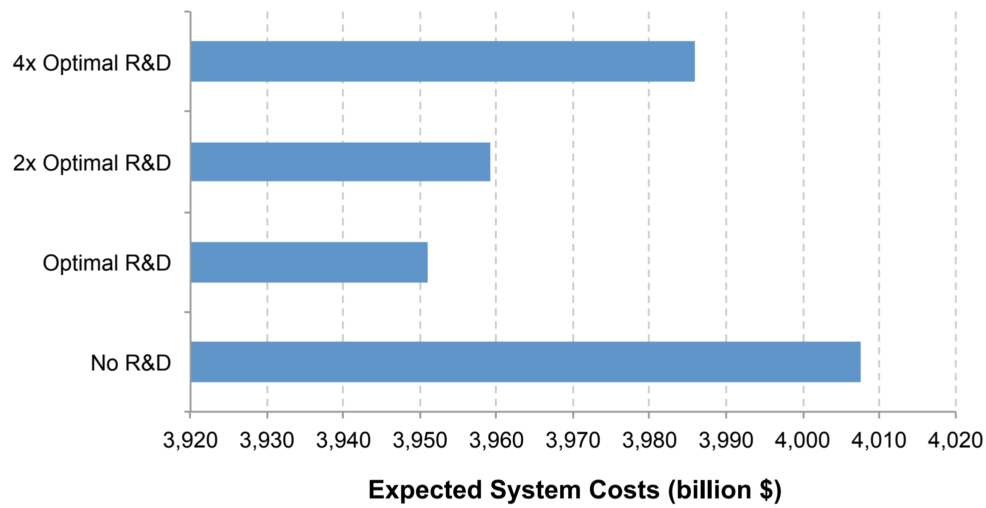

Figure 6.14: Comparison of expected system costs (billion $) under four R&D budgets.

There is an asymmetrical effect of overinvestment compared with underinvestment, as displayed in Figure 6.14. Overinvestment of a few times the optimal R&D budget has smaller cost increases relative to the case without any R&D investment, which echoes the results of Baker and Solak (2013). This asymmetry is caused largely by the assumption of decreasing marginal returns,[36] since the large productivity of initial R&D investments means that it would be more detrimental to make no expenditures than to make supererogatory R&D investments (up to a point). The degree of asymmetry depends on shapes of the innovation production functions and R&D

[36]If the R&D strategist thinks that R&D investments exhibit increasing returns, then the conclusion is flipped.

success values.

Figure 6.15 suggests that the net benefits from R&D equal zero when total investment across all programs is $7.63 billion, which is 5.6 times the optimal budget. The marginal benefits do not cross -$0.9 billion (per billion invested) until annual investment is $3 billion, which is 2.2 times the optimal budget.

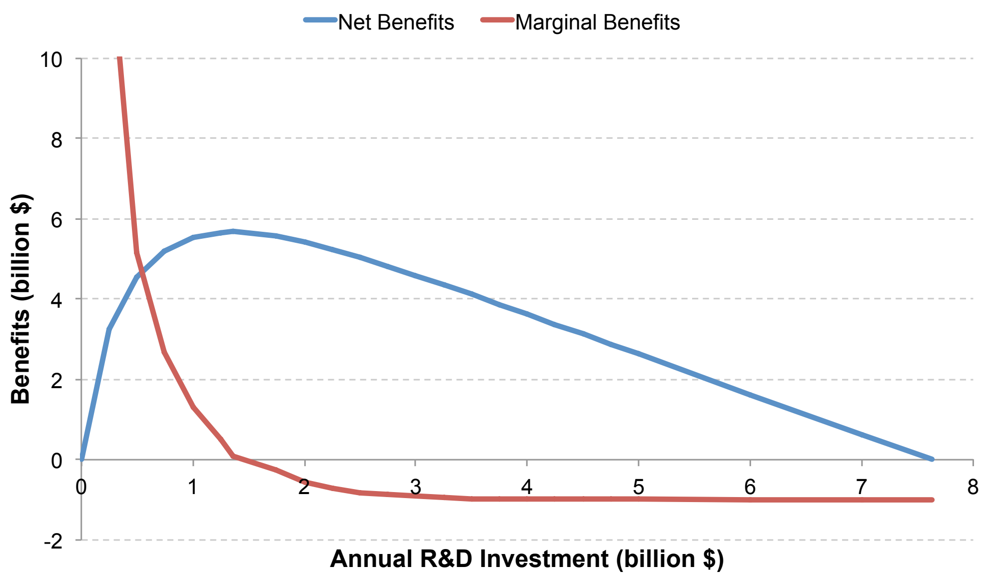

Figure 6.15: Net and marginal benefits for increasing R&D investments (billion $).

6.3.6 Nuclear Sensitivity

The results in previous sections indicate that nuclear R&D programs are attractive investments; however, these conclusions depend on the elicited probabilities for nuclear technologies. To test the robustness of these results, Figure 6.16 examines how the value of nuclear R&D success changes if different distributions are used. Given the wide variation of the expected value of R&D in a fairly small domain, more elicitations of how R&D success likely impacts costs for nuclear technologies should be conducted to verify the anticipated efficacy of the program.

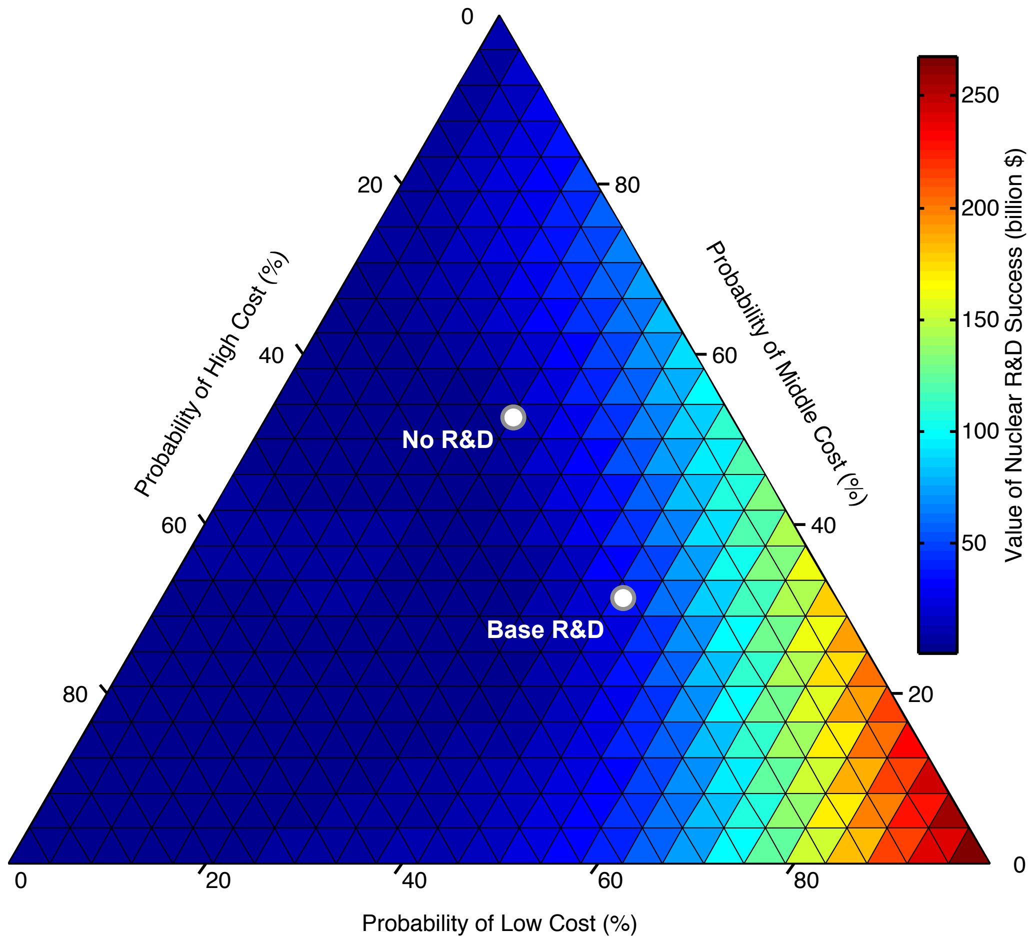

Figure 6.16: Probability simplex for the nuclear capital cost uncertainty. Gridlines represent different probabilities (which sum to one) for the cost outcomes of the three-point distribution. Colors on the simplex indicate the expected value of a nuclear R&D success given that the program moves the distribution over costs from the "No R&D" point to another point. The R&D success point used in previous sections is labeled "Base R&D."

As discussed in Chapter 4, investments in nuclear are economically attractive under a diverse range of market conditions due to its low life-cycle greenhouse gas emissions (which reduces exposure to the climate policy uncertainty) and to its low fuel price volatility (relative to alternatives like natural gas). However, the stochastic capacity planning model used in this research does not incorporate potential sociopolitical concerns or uncertainties surrounding nuclear power. These factors are inherently challenging to quantify due to the manifold concerns associated with these technologies (e.g., proliferation, waste disposal, safety) and to risk perceptions that are shaped by a complex combination of scientific risk assessments, cultural worldviews, and various psychological factors related to putative risk assessment (Kahan, 2012; Morgan et al., 2002; Douglas and Wildavsky, 1982). Incorporating these features would add additional risk to nuclear investments.

6.4 Summary and Extensions

This chapter presents a framework for directing investments in an energy technology R&D portfolio to promote innovation. A novel contribution of this work is to offer guidance about how R&D success valuations vary in different decision-making settings under uncertainty. Although theoretical results suggest that the effect of decision-making approaches on R&D valuation is equivocal, the numerical experiments using the two-stage stochastic capacity planning model indicate that R&D is more valuable in second-best planning environments like when decision-makers use expected-value approaches. These comparisons suggest that traditional, wait-and-see R&D valuation approaches likely overestimate the value of R&D success for many of the technologies but undervalue the optionality and hedging potential of technologies like nuclear relative to the stochastic approach.

The results stress the role of R&D in second-best policy environments. The expected value of R&D success is modulated by the degree to which environmental damages are internalized, and expectations about future policy decisions impact R&D investments in the present. In the presence of lax or nonexistent environmental policy, investments in R&D for more environmentally benign technologies are likely suboptimal. The results reinforce the common normative rationale that public R&D decisions should account for economic benefits accruing to all impacted stakeholders and not just to innovating firms.

A related conclusion is that R&D plays different roles under different policy environments, yielding the largest changes in environmental outcomes when abatement is low and the largest cost reductions when climate policies are stringent. This flexibility alludes to numerous real options associated with R&D budgeting, as a successful R&D strategy can assuage downside losses while enabling planners to capitalize on favorable investment opportunities. The results in this chapter illustrate the use of R&D as a method of reducing exposure to risk. Despite these dual roles, R&D alone is an inefficient means for achieving emissions reductions, which underscores the limitations of narrowly targeted policies for correcting environmental externalities and the significance of complementary policy instruments to incentivize mitigation across a range of timeframes. The environmental economics literature suggests that, if emissions prices prove not to be politically feasible, costs induced by political barriers will likely be substantial and that further constraints on available policy alternatives increase these costs (Fischer and Newell, 2008).

A broad goal of this framework is to assist decision-makers in thinking through the benefits of long-term, sustained R&D portfolio expenditures and to demonstrate the value of such programs. Although there are many exogenous uncertainties that may influence the diffusion of future technologies, judicious R&D spending requires the careful consideration of plausible scenarios under which these factors will align to allow for broad market penetration. The probabilistic modeling approach suggested in this chapter investigates a wide range of scenarios and identifies the conditions under which R&D successes will be most valuable (while also evaluating the likelihood of each scenario).

The proposed framework for R&D portfolio management offers a structure and set of tools for facilitating information gathering and model improvements. It is designed to offer support for difficult but necessary decisions in a way that makes its underlying dynamics and assumptions transparent but realistic. This framework can be updated as better information becomes available and as improved models are created. These traits are integral to the enterprise of managing technological change, as the characterization of innovation requires expert elicitations and other modeling judgments in areas where reasonable analysts and stakeholders may disagree. Although the model is applied here to a portfolio of energy technology R&D projects, the framework could be used to manage R&D in other domains like healthcare and national security.

The homogeneity of the optimal portfolio across a range of assumptions underscores both the importance of structure of the R&D valuation model and the expert elicitations. More detailed elicitations could reassess the validity of the innovation production functions used in the analysis and investigate whether some technologies exhibit increasing returns across an early investment range, which could be captured with a logistic function. The one-sidedness of the optimal portfolio should prioritize nuclear power elicitations with careful debiasing to avoid overconfidence. The recommendations of nuclear R&D support are based on relatively optimistic cost projections and should be tempered by the recognition of the historical overestimation of demand for nuclear power technologies and underestimation of their costs (Grübler, 2010; Hultman, Koomey, and Kammen, 2007; Koomey and Hultman, 2007; Cohen and Noll, 1991). On the other hand, this analysis points to the importance of explicitly considering market uncertainties when assessing R&D benefits, including the significance of nuclear as a hedging technology.

Potential extensions of this work include incorporating endogenous technical learning, quantifying spillovers across R&D programs, allowing R&D programs with variable time horizons, experimenting with a valuation model that includes other facets of the energy sector and a macroeconomic module, and modeling demand-side technology R&D programs.

Chapter 7

Fat-Tailed Uncertainty, Learning, and Climate Policy

"Not only does God definitely play dice, but He sometimes confuses us by throwing them where they can't be seen."

—Stephen Hawking

7.1 Background

Uncertainty is a pervasive feature in climate change economics. Deep structural uncertainty about the climate system combined with tremendous challenges in quantifying the economic impacts of severe climate change pose many conceptual, methodological, and ethical difficulties. In this context, the economic argument for more stringent near-term mitigation is typically based on reducing exposure to potential losses from catastrophic outcomes and not on mean-valued analyses from deterministic benefit-cost frameworks. Although insuring against low-probability, high-impact climate risks has been a leading justification for mitigation for many decades (Manne and Richels, 1993), recent developments have refocused attention on how fat-tailed uncertainty, where probabilities of rare events decline relatively slowly in the upper tail of a distribution,[1] may influence the urgency of abatement measures.[2]

Despite the centrality of risk in climate economics, the recognition of the importance of uncertainty has outpaced its actual quantification and implementation in the energy modeling community. Integrated assessment models (IAMs) are predominantly deterministic and focus on expected-value forecasts, using one-way sensitivities to assess the impact of uncertain or contentious parameters on model outputs. In the rare instances when uncertainty is more formally incorporated into the analysis, propagation techniques are typically used with thin-tailed probability density functions (Nordhaus, 2008; Hope, 2006).[3] These approaches offer a limited range of policy-relevant insights about near-term decisions and do not capture dynamics related to hedging, optionality, and learning, which come through the explicit inclusion of uncertainty in sequential decision-making frameworks (Kann and Weyant, 2000).

Weitzman (2009) highlighted these inadequacies regarding the treatment of uncertainty in IAMs and refocused research attention on the implications of extreme outcomes. Although the strong conclusions of the Dismal Theorem have been criticized on many fronts,[4] the paper has led to a reappraisal of fundamental notions of how to model (or, more generally, how to conceptualize) uncertainty, risk, discounting, and welfare in the context of fat-tailed environments with conceivably unlimited downside exposure. Despite these valuable contributions in framing the economic analysis of catastrophic risks and drawing attention to contentious assumptions in conventional models, the practical modeling implications and prescriptive policy guidance

[1]Fat-tailed distributions are defined here as having probabilities that decline polynomially or slower, as described in Section 7.2.2.

[2]The recent interests in the impacts and economics of low-probability, high-impact events, especially in relation to climate change, are reflected in Weitzman (2009); Frame and Allen (2008); Sunstein (2007); Taleb (2007); Weitzman (2007); Stern (2007); Parson (2007); Posner (2004).

[3]Popular choices for representing probability density functions are triangular (Hope, 2006) and normal (Nordhaus, 2008) distributions, which are frequently discretized. The *de-facto* practice of truncating distributions or imposing *de-minimis* risk thresholds, whether to simplify the computational complexity of the problem or to avoid assigning probabilities to difficult-to-quantify scenarios, potentially assumes away important dynamics of the climate system and the economics of climate change (Roe and Baker, 2007; Tomassini et al., 2007).

[4]The Dismal Theorem states that, "The catastrophe-insurance aspect of such a fat-tailed unlimited-exposure situation, which can never be fully learned away, can dominate the social-discounting aspect, the pure-risk aspect, and the consumption-smoothing aspect" (Weitzman, 2009). Millner (2013) summarizes and analyzes critiques of the Dismal Theorem.

for near-term decisions (e.g., how much insurance to buy) are left for future research.

In response, a growing body of literature in the economics of climate change has examined the impact of fat-tailed uncertainty on optimal abatement and the social cost of carbon. However, most of these studies investigate the implications of fat-tailed distributions using uncertainty propagation approaches instead of sequential decision-making ones. Studies like Dietz (2011); Pycroft et al. (2011); Ackerman, Stanton, and Bueno (2010); Mastrandrea and Schneider (2004); Tol (2003); Roughgarden and Schneider (1999) use Monte Carlo analyses to represent uncertainty by sampling from distributions for uncertain parameters, propagating them through a deterministic model, and creating output distributions. These approaches implicitly assume perfect information across the entire time horizon for each simulation run (i.e., implying a learn-then-act approach in which the uncertain state is revealed before decisions are made). This characteristic means that such *ex-post* approaches, while analytically simple, cannot offer guidance in determining *ex-ante* hedging strategies, which balance the risks of premature action with those of delay given the decision-maker's present state of knowledge. Sequential decision-making frameworks incorporate uncertainty explicitly and address limitations of foresight by determining optimal policies in multiple stages based on updated information.[5] These adaptive approaches more realistically capture the decision environment for climate policy, where perfect information is not likely to be available in the immediate future.

The objective of this research is to examine how sequential decision-making frameworks, concepts, and metrics can be used to inform risk management in climate policy. Specifically, this work examines the impact of fat-tailed uncertainty about the climate sensitivity parameter and of the potential for learning on optimal near-term abatement.[6] Model experiments answer questions about whether policy recommendations

[5]Weitzman (2012) implicitly suggests that climate risk should be accounted for in a sequential decision-making framework rather than through uncertainty propagation, since the Monte Carlo treatments of uncertainty "very likely fail to account adequately for the implications of large impacts with small probabilities."

[6]For illustrative purposes, this analysis focuses on uncertainty about the climate sensitivity parameter as an aggregate measure for the climate response to greenhouse gas emissions, which reflects the parameter's prominent role in many IAMs and its correlation with climate change effects (Pycroft et al., 2011). However, the analysis also could introduce uncertainty explicitly for economic impacts (where deep uncertainty exists) or for climate-related thresholds and feedbacks.

are robust to the specification of distributions, damages, and discounting.

Unlike Weitzman (2012), this work offers prescriptive policy guidance under fat-tailed uncertainty beyond the restrictive Dismal Theorem. Criticisms of the limiting assumptions of the Dismal Theorem (Nordhaus, 2012) suggest that fat tails *per se* are not sufficient to lead to an unbounded expected utility but do not provide guidance for policy selection under fat-tailed uncertainty more generally. Even if the strong Dismal Theorem conditions proposed by Weitzman do not obtain, the tail obesity of the climate sensitivity parameter may be appreciable enough to merit consideration outside of the standard deterministic benefit-cost setting. The modeling framework in this chapter makes amendments to deterministic IAMs to avoid this policy impasse by incorporating uncertainty through a sequential decision-making framework.[7]

Another goal of this research is to quantify the value of learning and midcourse corrections on reducing consumption risks imposed by uncertain damages from climate change. In the presence of strong stock-accumulation inertias and sunk emissions, Weitzman (2012) posits that the possibility of learning is irrelevant and that it would be challenging to change course if the response of the climate system is more severe than expected.[8] Thus, Weitzman (2012) assumes that lags involved in climate system preclude the ability to learn about catastrophic impacts until damages arrive in 150 years. Other authors (Nordhaus, 2012; Kousky et al., 2010; Nordhaus, 2009; Yohe and Tol, 2007) have criticized Weitzman's assumptions as being unrealistically pessimistic and suggest that the pace of climate change will allow for the possibility of midcourse corrections over time, of deploying negative emissions technologies, or of using emergency geoengineering if warming is unexpectedly high.[9] Despite these

[7]The analysis strikes a balance between rigor and practical utility by using a simplified stochastic IAM, which incorporates dynamics of sequential decisions while adopting more stylized climate and investment dynamics, as discussed in Section 7.2. It offers policy guidance in terms of optimal abatement instead of willingness to pay (Weitzman, 2012; Pindyck, 2012).

[8]Weitzman (2012) emphasizes this point by citing the "crude rule of thumb" from Solomon et al. (2009) that about 70 percent of the peak CO_2 enhancement level will persist after a century of no emissions and 40 percent will remain in the atmosphere after a millennium.

[9]Although the learning rate is uncertain and dependent on factors like the true value of the climate sensitivity parameter (Kelly and Tan, 2013), the high likelihood that information will be available early enough to enable midcourse corrections makes the economic analysis of climate change amenable to and well-suited for a multi-stage stochastic programming with recourse framework, which is used here. In this respect, climate change may offer greater potential for learning and

criticisms, no work has investigated the degree to which assumptions about learning influence the conclusions of the model in Weitzman (2012). This research investigates the influence of different learning rates on near-term abatement.

Although previous work has investigated the learning process associated with climate change (Webster, Jakobovits, and Norton, 2008; Leach, 2007; Yohe, Andronova, and Schlesinger, 2004; Webster, 2002; Kelly and Kolstad, 1999), nearly all of these analyses have not explored such issues in the presence of fat-tailed uncertainty. The limited research conducted on fat-tailed uncertainty using sequential decision-making frameworks has not compared learning in the context of different priors and often focuses on the rate of learning. In general, such analyses (Kelly and Tan, 2013; Gerst, Howarth, and Borsuk, 2010) do not simultaneously test the specifications of uncertainty, learning, damages, and discounting in a unified framework and make restrictive assumptions, which limit the generalizability of the results.

A high-level motivation for this work is to investigate to what degree near-term policy prescriptions in IAMs are robust to conventional assumptions about thin-tailed probabilities, perfect foresight, and quadratic damages. The following sections describe the deliberately parsimonious model and experiments that are designed to isolate these effects and to show how policy guidance from IAMs may rely strongly on these assumptions in certain domains. The results do not propose specific and definitive solutions to these complex issues but draw general insights, which are useful starting points in understanding modeling limitations and refocusing research attention on problems related to uncertainty and learning.

7.2 Model

The model formulation is based on a simplified version of the Dynamic Integrated Climate-Economy (DICE) model (Nordhaus, 2008) with extensions to account for

midcourse corrections relative to other types of catastrophes (e.g., catalyzed conversion by strangelets from heavy-ion collisions in high-energy particle accelerators).

uncertainty through a two-stage sequential decision-making framework.[10] For comparability, it makes similar parameterization assumptions to the Weitzman (2012) model. The social planner's objective is to determine the optimal consumption path for goods and services and for investments in technologies to reduce greenhouse gas emissions, which balances the cost of mitigation and damages from climate change.

7.2.1 Economic and Climate Systems

The model finds the optimal path of abatement and consumption to maximize the expected net present value of discounted utility flows:

$$\max_{\mu_1} \int_0^{\tau} e^{-\rho t} U(C_t, L_t)\, dt + \mathbb{E}\left[\max_{\mu_2} \int_{\tau}^{\infty} e^{-\rho t} U(C_t, L_t)\, dt\right] \tag{7.1}$$

The choice of the first-stage control rate (μ_1) depends on the uncertain and unobservable value of the climate sensitivity parameter, which is not known until time τ. The control rate represents the aggregate percentage reduction of greenhouse gas emissions below business-as-usual levels across the decision stage.[11] The nonanticipativity assumption constrains first-stage decisions to be identical under all possible states of the world, which makes the optimal control choice a sequential decision under uncertainty.[12]

The population of L_t identical agents has preferences over per capita consumption

[10]The notation in this chapter is independent of other dissertation chapters, so previously used symbols may have different meanings in this model.

[11]This control rate is the average value across the duration of the first stage and not the immediate policy value, as there are many abatement paths consistent with this average rate.

[12]Hwang, Reynès, and Tol (2011) is one of the only papers claiming to use stochastic programming to examine fat-tailed climate uncertainty. However, since consumption and abatement decision variables in their model are selected for each state of the world without an additional nonanticipativity constraint, their mathematical formulation is functionally equivalent to using a wait-and-see approach. This implicit assumption of perfect information about the true scenario before first-stage decisions are selected makes it indistinguishable from the expected value of the wait-and-see approach. Without the possibility of making suboptimal decisions during first stage, this approach underestimates the impact of uncertainty on optimal policy selection.

described by the isoelastic (constant relative risk aversion) utility function:

$$U(C_t, L_t) = L_t \frac{\left(\frac{C_t}{L_t}\right)^{1-\eta} - 1}{1-\eta} \tag{7.2}$$

The labor supply L_t is assumed to be inelastic and equal to the population. The population grows exogenously at a rate n until reaching a steady-state asymptotic maximum L_∞.

The optimization problem is subject to the constraint that net output[13] (Y_t) after damages ($\widetilde{\Omega}$) at time t is equal to the sum of consumption, capital investment, and abatement costs:

$$\widetilde{\Omega} Y_t = C_t + I_t + \Lambda_t \tag{7.3}$$

Here, the savings rate (σ) is assumed to be an exogenously specified fraction of net output in time t. Additionally, output grows at an exogenous rate g:

$$\frac{\dot{Y}_t}{Y_t} = g \tag{7.4}$$

This growth rate can be conceptualized as total factor productivity and is a combination of capital- and labor-augmenting technical progress.

The convex abatement cost function is defined by the equation:

$$\Lambda_t = B_t \mu_i^\theta Y_t \tag{7.5}$$

Economic output is adjusted by the economic loss (i.e., damage) function given by the formula:

$$\widetilde{\Omega} = \left[1 + \left(\frac{T}{\alpha}\right)^2 + \left(\frac{T}{\beta}\right)^\gamma\right]^{-1} \tag{7.6}$$

where T represents the equilibrium change in the global mean surface temperature, and $\widetilde{\Omega}$ is expressed as a fraction of global economic output after accounting for losses due to climate damages. For larger values of T, climate change negatively impacts

[13]Net of capital depreciation.

output in the second stage through this multiplicative factor. Although atmospheric equilibrium may take many centuries for large temperature changes to materialize, damages are conceptualized as being linked to the trajectory with this asymptotic limiting temperature change and are normalized as if they occur after τ_D years. This analysis makes the strong assumption that all climate change impacts (i.e., market and non-market) can be monetized and expressed as a fraction of output.

Given criticisms that quadratic-polynomial damage functions likely understate welfare impacts for large temperature changes, the results explore impacts of using a reactive damage function with a parameterization that accounts for the possibility of more appreciable economic impacts at higher warming levels (i.e., economic damages are more reactive to higher temperature changes than a quadratic damage function).[14] As discussed in Section 7.2.3, the calibrated values for the reactive damage function come from Weitzman (2012). In Equation 7.6, the rightmost term is equal to zero for quadratic damages and has an exponent of $\gamma = 6.754$ for reactive damages.

The model makes the simplistic assumption that the greenhouse gas concentration builds to a level G (conditional on abatement choices) in τ_D years when damages are suddenly realized. Output is reduced by the fraction $\widetilde{\Omega}$ associated with the realization of the equilibrium temperature conditional on the temporally aggregated stock of greenhouse gas emissions in the atmosphere.

The forcing factor as a function of the steady-state atmospheric carbon dioxide (CO_2) concentration G (in parts per million by volume) is:

$$\Phi = \frac{\ln(G/280)}{\ln(2)} \tag{7.7}$$

where the function is normalized so that, at the atmospheric doubling concentration from the pre-industrial level, $\Phi(G = 560) \equiv 1$. For analytical convenience, the model uses a simplified mapping between the emissions profile and atmospheric stock

[14] As pointed out in Weitzman (2012), quadratic damage functions like the one found in DICE are based on estimates of economic impacts from small temperature changes, which makes extrapolation questionable for higher temperatures. For instance, the parameterization in DICE leads an 8 percent reduction in output for a temperature increase of 6 °C and a 44 percent loss for a temperature increase of 18 °C.

of greenhouse gas emissions.[15] The steady-state concentration depends on the CO_2 retention rate in the atmosphere (ϕ) and first- and second-stage emissions (μ_1 and μ_2, respectively) according to the equation:

$$G = \phi G_0 + (1-\mu_1)\frac{\tau}{\tau_D}\Gamma + (1-\mu_2)\frac{\tau_D - \tau}{\tau_D}\Gamma \tag{7.8}$$

where Γ is the uncontrolled CO_2 concentration increase after τ_D years from DICE-2010. The steady-state temperature change as a function of the atmospheric CO_2 concentration and the equilibrium climate sensitivity parameter S is given by $T = \Phi(G) \times S$. In an extremely simplified sense, global temperature changes from anthropogenically injected CO_2 are approximately the product of emissions and an uncertain climate-sensitivity-like scaling parameter, which is discussed in the following section.

7.2.2 Climate Sensitivity Parameter Distributions

The driving random variable in this analysis is the equilibrium climate sensitivity parameter (S). The climate sensitivity parameter is defined as the global mean surface temperature increase resulting from a sustained doubling of atmospheric greenhouse gas concentrations after the climate has reached a new steady-steady equilibrium.[16] This parameter is an aggregate indicator of the climatic response to stocks of anthropogenically injected greenhouse gases.

Like Weitzman (2012), this work treats the climate sensitivity parameter as a reduced-form representation of the many uncertain dimensions of climate change, since it is heavily researched and offers greater specificity to the problem. In actuality, even if the true value of the climate sensitivity were known, there would still be substantial uncertainty about the response of the Earth system and about spatial and temporal distributions of localized impacts. Properly treating the uncertainty

[15] Although these equations offer an enormous simplification of the climate system, the model is similar to other representations used in climate policy analysis. More complex equations of motion likely would not alter the qualitative conclusions from this simplified model.

[16] For this analysis, the climate sensitivity parameter refers to the fast equilibrium value, which does not incorporate the slower S_2 feedbacks (e.g., changes in biological sources or sinks, temperature-induced greenhouse gas releases, and others) considered in Weitzman (2009).

associated with a coarse parameter like the climate sensitivity (though critical in determining an aggregate policy response) does not tell us much about when and where climate change impacts will be felt, which leaves many residual economic, policy, and ethical questions.

Table 7.1: Alternate probability distributions for the climate sensitivity parameter.

Distribution	**Density Function**	**Support**	**Parameters**
Normal	$f_{\mathcal{N}}(S;\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(S-\mu)^2}{2\sigma^2}}$	$S \in \mathbb{R}$	$\mu = 3$ $\sigma = 1.447$
Lognormal	$f_{\mathcal{L}}(S;\mu,\sigma) = \frac{1}{S\sigma\sqrt{2\pi}} e^{-\frac{(\ln S-\mu)^2}{2\sigma^2}}$	$S \in (0,+\infty)$	$\mu = 1.099$ $\sigma = 0.3912$
Pareto	$f_{\mathcal{P}}(S;\pi,\alpha) = \frac{\alpha\pi^\alpha}{S^{\alpha+1}}$	$S \in [\pi,+\infty)$	$\pi = 2.3758$ $\alpha = 2.969$

The model uses three different distributions based on Weitzman (2012) to represent various levels of tail obesity, as shown in Table 7.1. *Thin-tailed distributions* have probabilities that decline exponentially or faster, which is represented in this research by the normal distribution. Probabilities in the upper tails of *intermediate-tailed distributions* decline slower than exponentially but faster than polynomially, which is embodied in the lognormal distribution. *Fat-tailed distributions* are defined here as having probabilities that decline polynomially or slower, which is represented by the Pareto distribution.[17]

The Intergovernmental Panel on Climate Change's Fourth Assessment Report states that:

> [The equilibrium climate sensitivity] is *likely* to be in the range 2 °C to 4.5 °C with a best estimate of about 3 °C, and is *very unlikely* to be less than 1.5 °C. Values substantially higher than 4.5 °C cannot be excluded, but agreement of models with observations is not as good for those values.

[17] There are many definitions of tail obesity (Cooke, Nieboer, and Misiewicz, 2011), which reflects the long history of fat tails and leptokurtic distributions in mathematics (Mandelbrot, 1963).

Although the IPCC does not attach a probability to the climate sensitivity being greater than 4.5 °C ($P[S > 4.5\ °C]$), the IPCC defines "likely" elsewhere as a probability above 66 percent but below 90 percent.[18] Weitzman (2012) calibrates the three distributions so that $P[S > 3\ °C] = 0.5$ and $P[S > 4.5\ °C] = 0.15$.

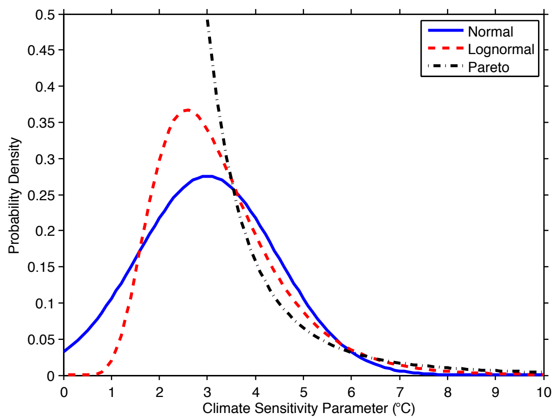

(a) Distributions

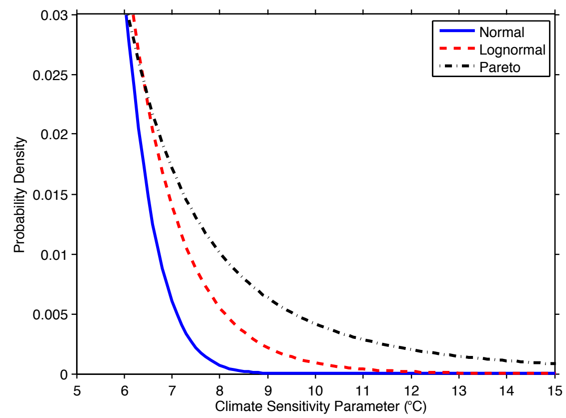

(b) Close-ups of upper tails

Figure 7.1: Alternate probability density functions for the climate sensitivity parameter with various levels of tail obesity.

Although the probabilities of very high climate sensitivity values are small, Figure 7.1 shows that the relative tail probabilities depend greatly on the specification.[19] Thin-tailed distributions essentially exclude the possibility of climate sensitivity values greater than about 8 °C.

[18] Kunreuther et al. (2012) note that estimates of $P[S > 4.5\ °C]$ in the literature range from less than 2 percent to as high as 50 percent. Elicitation results from Zickfield et al. (2010) suggest that this probability is approximately 23 percent.

[19] For instance, $P[S > 10\ °C]$ is 1.4 percent for the Pareto distribution, 0.1 percent for the lognormal, and 7×10^{-7} percent for the normal.

7.2.3 Calibration

The model adopts a descriptivist approach for selecting the parameters that influence the discount rate in the Ramsey equation:

$$r_t = \rho + \eta g_t \tag{7.9}$$

The values for the pure rate of time preference (ρ) and elasticity of marginal utility (η) are inferred from decisions made by democratically elected governments using mortality rates and personal income tax structures, respectively. These revealed social values are based on the mean values from Anthoff, Tol, and Yohe (2009) using data from Evans and Sezer (2005, 2004). The parameter values of $\rho = 1.08$ percent and $\eta = 1.49$, combined with growth rate of $g = 2.1$ percent (La Grandville, 2012), give a discount factor that aligns with a 5 percent rate of time discount.[20]

The damage function in Equation 7.6 can represent either quadratic or reactive damages depending on the chosen parameter values. For quadratic damages, the rightmost term drops out as $\gamma \to -\infty$, and the equation becomes identical to the damage function found in DICE (Nordhaus, 2008). For reactive damages, the α parameter is the same, and $\beta = 6.081$ and $\gamma = 6.754$ are selected to match Weitzman (2012), which calibrates the damage function so that $\Omega(T = 6\ °\mathrm{C}) = 0.5$. Similarly, the analysis here adopts Weitzman's related assumptions that economic damages from climate change do not materialize for $\tau_D = 150$ years and that 18 °C is taken as an upper bound on the temperature change from preindustrial levels.[21] *Ceteris paribus*, these assumptions bias mitigation downward.

The marginal abatement cost curve is calibrated so that optimal first-stage abatement without learning is 0.27 using a normal distribution and quadratic damages, which is the average control rate until 2105 in DICE-2007.

[20]The secondary motivation for this choice is to select a comparatively high discount rate so that model results cannot be attributed primarily to an unreasonably low discount rate. Ultimately, the descriptivist parameters may not reflect normative social preferences (Kaplow, Moyer, and Weisbach, 2010). Such differences may be critical for policies with long timeframes and deep uncertainty.

[21]Costello et al. (2010) argue that the equilibrium temperature change should be truncated due to the physical impossibility of infinite temperatures, which is implemented here with the same threshold from Weitzman (2012).

The exogenous savings rate comes from DICE-2010. This rate as a fraction of net output does not exhibit temporal variation in DICE-2010 and is nearly identical across different abatement paths.[22]

The population growth rate ($n = 0.01$) is from La Grandville (2012), and the asymptotic population ($L_\infty = 8.7$ billion) comes from DICE-2010.

Table 7.2: Calibrated parameter values for the fat-tailed uncertainty model.

Symbol	Parameter	Value
τ	learning time (years)	150
τ_D	impact time (years)	150
ρ	pure rate of time preference	0.0108
η	elasticity of marginal utility	1.49
L_0	initial population (million)	6,411
L_∞	asymptotic population (million)	8,700
n	labor force growth rate	0.01
ω	savings rate	0.225
g	output growth rate	0.021
B_τ	cost of backstop abatement technology (\$/Mt-C)	1,000
ϕ	atmospheric CO_2 retention rate	0.877
G_0	atmospheric CO_2 concentration at $t = 0$ (ppmv)	389
Γ	uncontrolled CO_2 concentration increase by 2155 (ppmv)	699

7.3 Results

7.3.1 Reference Results

In this section, it is assumed that the social planner does not learn the true equilibrium climate sensitivity parameter until 2150, which corresponds to the reference assumptions of Weitzman (2012). The model used here calculates the optimal near-term abatement policy as the output of interest (instead of willingness to pay to avoid climate change).

[22]The time invariance of the propensity to consume is similar to Gerst, Howarth, and Borsuk (2010). This paper cites data from the International Monetary Fund suggesting that the savings rate has been stable in recent decades.

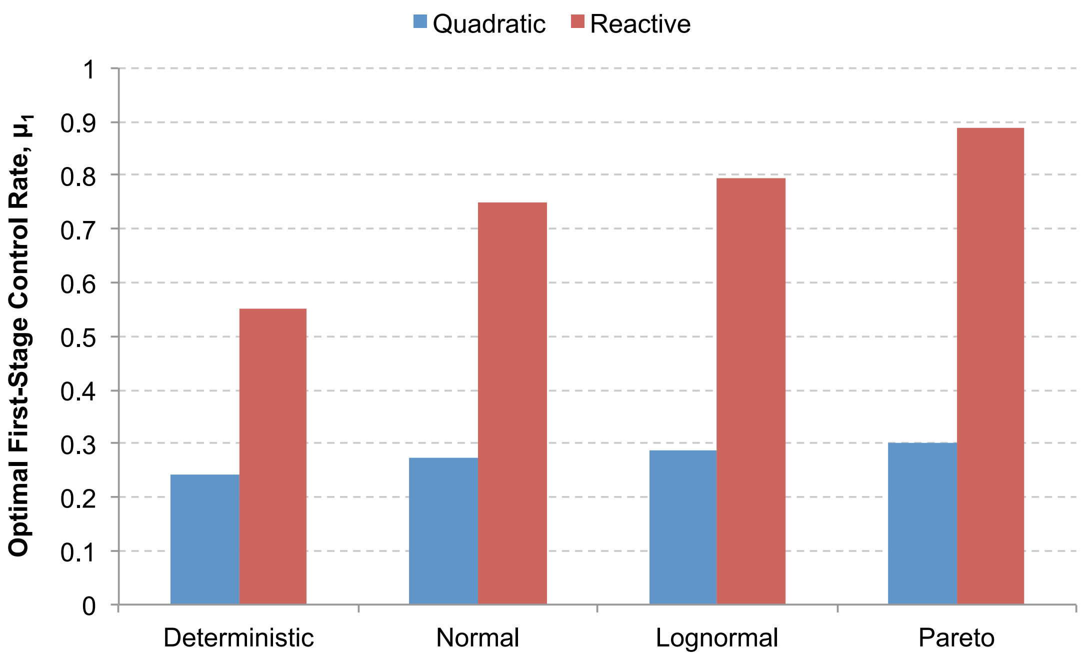

Figure 7.2: Comparison of optimal first-stage control rates (μ_1) under alternate climate sensitivity parameter distributions and damage functions when $\tau = 150$.

A motivation for this research is to ascertain the degree to which alternate distributions for the climate sensitivity have different implications for optimal abatement. Figure 7.2 shows that the answer to this question depends strongly on the assumed damage function. For quadratic damages calibrated to estimates for modest temperature changes (Nordhaus, 2008), tail thickness does not play a large role in determining near-term policy due to relatively small consumption losses even at high temperatures. This result suggests that questions of abatement under quadratic damages belong to a class of problems with "tail irrelevance," where the "distribution of the random variable makes no (or little) difference to the policy or the outcome" (Nordhaus, 2012). This result also gives a sense for why uncertainty analyses with quadratic damages often conclude that uncertainty does not matter. For instance, using DICE in an uncertainty propagation analysis with quadratic damages and normal distributions, Nordhaus (2008) concludes:

> Look at the current (2005) social cost of carbon, we see that the mean

> estimate ($26.85 per ton) is slightly less than the most likely estimate ($28.10 per ton). This important finding indicates that the estimates in the certainty-equivalent model are very close to the estimates in the uncertainty model.

In contrast, fatter-tailed uncertainty increases abatement relative to the thin-tailed and deterministic cases when reactive damages are assumed.[23] Fat-tailed uncertainty without learning increases the control rate by 61 percent relative to the deterministic case and 18 percent relative to the normally distributed case. Thus, the results substantiate the non-robust dependence of the recommended policy action on the specification of uncertainty (specifically, to the assumed thickness of the upper tail) and damages. Even within a benefit-cost framework, the results offer similar qualitative conclusions to Weitzman (2012) when the possibility of learning is excluded (i.e., $\tau = \tau_D = 150$ years).

To place these first-stage emissions control rates in perspective, the average control rate for the deterministic DICE-2010 model between 2005 and 2155 is 0.55. This value falls between the deterministic and normally distributed control rates with reactive damages in Figure 7.2. When the DICE model is used with the discounting parameters suggested by the *Stern Review* (Nordhaus, 2008), the average control rate over this same period is 0.83, which is comparable to the value from the case with reactive damages and fat-tailed Pareto uncertainty (0.89). Thus, one reason why Weitzman suggests that the policy conclusions from the *Stern Review* may be "right for the wrong reasons" (Weitzman, 2007) is that the recommended climate policy stringency is similar between the low-discounting, thin-tailed uncertainty and moderate-discounting, fat-tailed uncertainty cases. However, this conclusion only holds when reactive damages are assumed.

[23]The deterministic case assumes that the climate sensitivity parameter is 3 °C with a probability of one.

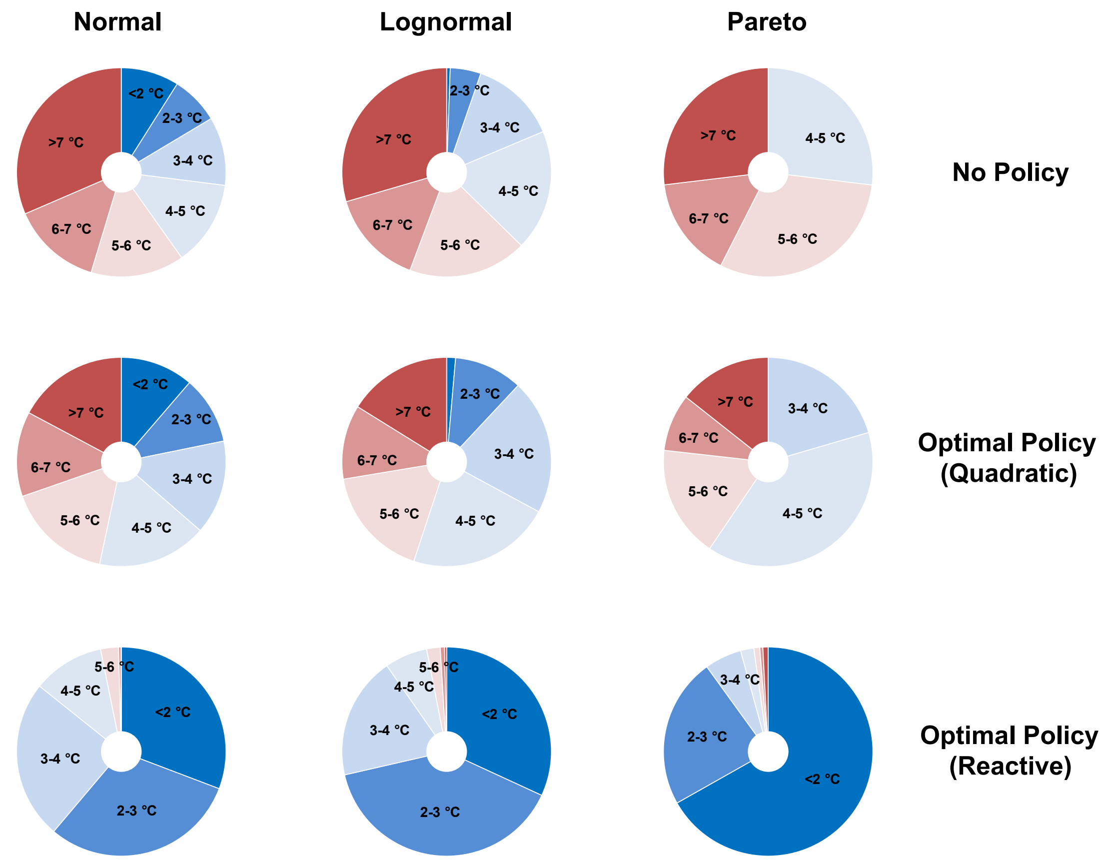

Figure 7.3: Probabilities of equilibrium global surface temperature changes (°C) under different climate policies and distributions over the climate sensitivity parameter. The size of each slice reflects the likelihood of a specified warming outcome.

Figure 7.3 illustrates how climate change risk is strongly influenced by the target mitigation level and also by the assumed distribution for the climate sensitivity parameter. This figure uses roulette wheel visualizations (Webster et al., 2012) to convey the uncertainty associated with future temperature changes and the influence of abatement decisions on such outcomes. Under the no-policy scenario with a normally distributed climate sensitivity, there is a 31 percent chance that the global mean surface temperature will increase by at least 7 °C (12.6 °F) and a 16 percent chance that warming will be limited to less than 3 °C (5.4 °F). For the other distributions, the no-policy scenario leads to similar probabilities of high temperature increases, but the fatter-tailed uncertainties lead to smaller probabilities of warming below 4 °C.

The policy cases in the lower two rows show how even comparatively small policy changes can trim the probabilities associated with large temperature changes compared with the inaction scenario. Under the climate policy recommended with quadratic damages, the probabilities associated with significant levels of warming are reduced for all three distributions. For instance, the probability that temperature increases are lower than 5 °C is greater than 50 percent for all three distributions. However, there is still an appreciable probability of high temperature increases when only modest mitigation efforts are undertaken. In contrast, climate policies suggested under reactive damages illustrate how policy can attenuate deep structural risks, as the probability of temperature increases above 5 °C are under five percent. Thus, Figure 7.3 conveys how mitigation can function as insurance against economic and welfare risks imposed by climate change.

7.3.2 Value of the Stochastic Solution

Another motivation for this work is to examine how sequential decision-making frameworks and metrics, like those developed and used elsewhere in this dissertation, can inform climate policy. As described in Chapter 3.3.3, the value of the stochastic solution (VSS) compares the expected utility of the optimal stochastic solution with one that ignores uncertainty by assuming that the expected-value of a distribution will be realized and then solving the resulting deterministic problem. The VSS quantifies the value of explicitly including uncertainty in planning efforts.

Table 7.3 suggests that analyzing uncertainty is important in modeling efforts and policy design when economic damages associated with temperature increases are more nonlinear and when the distribution over the climate sensitivity parameter is asymmetrical.[24] As indicated in Weitzman (2012), it is the combination of unknown climate outcomes and the potential for more severe damages in those states that make the inclusion of uncertainty important in decision-making.

[24]Greater emphasis should be placed on relative changes and not on absolute values associated with the objective function.

Table 7.3: Objective function value and VSS (trillion $) under various assumptions about climate sensitivity parameter distributions and damage functions.

	Quadratic			Reactive		
	z^*	z_d	**VSS**	z^*	z_d	**VSS**
Normal	1,696	1,696	0	1,686	1,676	10
Lognormal	1,695	1,695	0	1,683	1,666	17
Pareto	1,695	1,695	0	1,681	1,629	52

7.3.3 Learning

The primary objective of this work is to determine the value of learning and mid-course corrections on reducing consumption risks imposed by uncertain damages from climate change. This section answers the question of when uncertainty needs to be resolved to make near-term policies differ from the no-learning scenario described earlier. These experiments illustrate the effects of learning on policy choice using three possibilities for learning:

- *No learning*: The true value of the climate sensitivity parameter is not known for 150 years ($\tau = \tau_D$), which is assumed under the reference conditions in earlier sections.

- *Early learning*: Perfect information about the true value of the climate sensitivity parameter is available before damages arrive ($\tau < \tau_D = 150$), which allows mitigation decisions to be revised in light of this early learning.

- *Perfect information*: The perfect information case refers to the scenario where the decision-maker has immediate access to the true value of the climate sensitivity parameter ($\tau = 0$).

Although no case is entirely realistic, these experiments bound the range of abatement across possible assumptions about learning.

Figure 7.4 graphs the first-stage control rate as a function of date when uncertainty is resolved. The expectation of learning has a significant impact on the optimal near-term policy. Perfect information at the beginning of the time horizon reduces first-stage abatement effort by 41–48 percent (depending on the assumed prior) relative to the condition of uncertainty but no learning.[25] This result suggests that Weitzman's conclusions (i.e., the non-robustness of policy recommendations to the representation of uncertainty in IAMs) are non-robust to the model specification of learning.

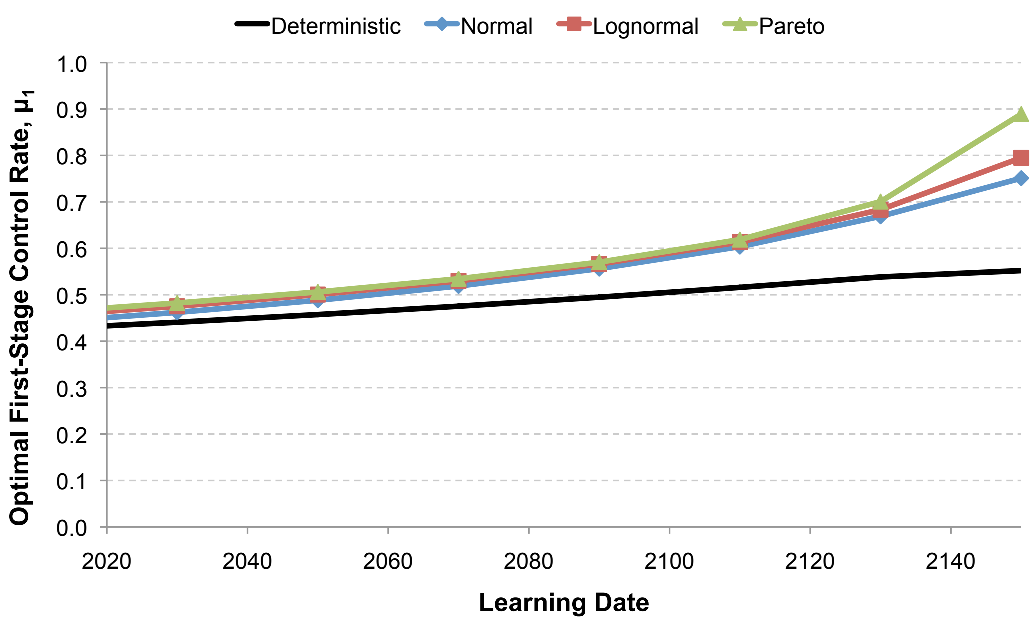

Figure 7.4: Optimal first-stage control rates (μ_1) under reactive damages and alternate climate sensitivity parameter distributions as a function of the learning date.

Since learning can affect first-stage abatement, these results suggest that IAMs do not need frictions or learning-by-doing to induce anticipatory actions that differ from the expected-value solution. Here, learning provides the capacity to reduce regret associated with early-stage decisions, which is especially beneficial when the true value of a random variable differs from the expected value. For instance, assuming a lognormal distribution and reactive damages, the average control rate through 2150 when

[25]These magnitudes are similar to the results from Kelly and Tan (2013).

learning occurs in next 80 years is 57 percent, which is similar to the literature (e.g., the DICE-2010 optimal control rate is about 55 percent). However, without possibility of learning, the control rate is approximately 80 percent. In the absence of learning, the planner must insure more during initial stages given that the possibility of making midcourse adjustments is excluded. The combination of fat-tailed uncertainty, reactive damages, and lengthy learning times strengthens inter-stage dependencies by making second-stage damages more strongly influenced by earlier abatement decisions. In this way, learning can be thought of as reducing the risk premium and the exigency of precautionary insurance, even though the results demonstrate that at least some precautionary abatement is warranted across a wide range of assumptions about uncertainty, learning, damages, and discounting.

A second takeaway from Figure 7.4 is that the choice of climate sensitivity distribution, provided that uncertainty is incorporated in some form, does not have a sizable impact when learning is expected in the next century. If information arrives in a sufficiently short time, the sunk emissions during early periods do not lock in commitments to warming, and the decision-maker can still adjust emissions flows to ensure that the atmospheric stock of greenhouse gases does not reach threshold levels associated with large losses. If learning does not occur in this timeframe, irreversible stock accumulations make adjustments less effective, and fat-tailed distributions attach increasingly smaller probabilities to minor temperature increases. Although the actual learning rate of the climate sensitivity will depend on a host of factors including the true value of the parameter (e.g., since larger values will lead to longer learning times), Kelly and Tan (2013) suggest that the expected time until complete learning is about 80 years.[26] Ultimately, the policy-relevance of potential fat tails depends on beliefs about learning and damages.

Finally, learning is relevant no matter what level of tail obesity a planner assumes as long as uncertainty is incorporated explicitly into the decision problem. This result is a testament to the policy importance of uncertainty in general and learning in particular. Figure 7.4 illustrates how wait-and-see (i.e., perfect information) approaches

[26] Kelly and Tan (2013) show that complete learning may take up to two centuries if the true value is in the tails, but the potential learning times to rule out fat tails are between 12 and 39 years.

(Hwang, Reynès, and Tol, 2011) recommend control rates and carbon taxes that are likely biased downward.

Table 7.4 helps to identify cases where learning impacts near-term policy decisions. The expectation of early learning only has a significant impact on the optimal first-stage control rate when the damage function is reactive. With reactive damages, the expected value of early information is high, and if perfect information could be available in 2100, the value of information is \$71 and \$124 trillion for the normal and Pareto distributions, respectively. Learning does not have an appreciable abatement impact when the pure rate of time preference is very low, since low values induce high abatement regardless of the learning rate.

Table 7.4: Optimal first-stage control rates (μ_1) as a function of learning under alternate assumptions about the climate sensitivity parameter distribution, damage function, and pure rate of time preference (ρ). The distribution cases assume reactive damages, and the damage function cases assume a lognormal distribution. The baseline ρ is 1.08 percent.

	Distribution			**Damages**		ρ	
	$\mathcal{N}$	$\mathcal{L}$	$\mathcal{P}$	**Quadratic**	**Reactive**	**0%**	**1.08%**
No Learning	0.75	0.80	0.89	0.29	0.80	0.90	0.80
2100	0.58	0.59	0.59	0.26	0.59	0.84	0.59
2050	0.49	0.50	0.51	0.24	0.50	0.83	0.50
Perfect Information	0.44	0.45	0.46	0.23	0.45	0.83	0.45

7.3.4 Elasticity of Marginal Utility and Risk Aversion

The isoelastic utility function does not differentiate between risk aversion, aversion to intertemporal inequality, and aversion to intratemporal inequality, which are all controlled by the elasticity of marginal utility parameter (η). However, psychological research suggests that individuals are not equally averse to risk and inequality (Atkinson et al., 2009; Traeger, 2009). Thus, an appropriate value for the societal η parameter is unclear.

Figure 7.5 compares first-stage control rates for various assumptions about η when $\tau = 150$ and reactive damages are assumed. The U-shaped relationship reflects the dual role of the parameter both as the elasticity of marginal utility of consumption and as the coefficient of relative risk aversion. The first effect tends to lower the control rate as η increases (due to discounting), and the second effect tends to increase the control rate (due to risk aversion). The relationship between optimal abatement and changes in risk preference in Figure 7.5 shows that the diminishing marginal utility effect is stronger than risk aversion in this context. The minimum control rate value depends on assumed climate sensitivity distribution, though the general trend is similar.

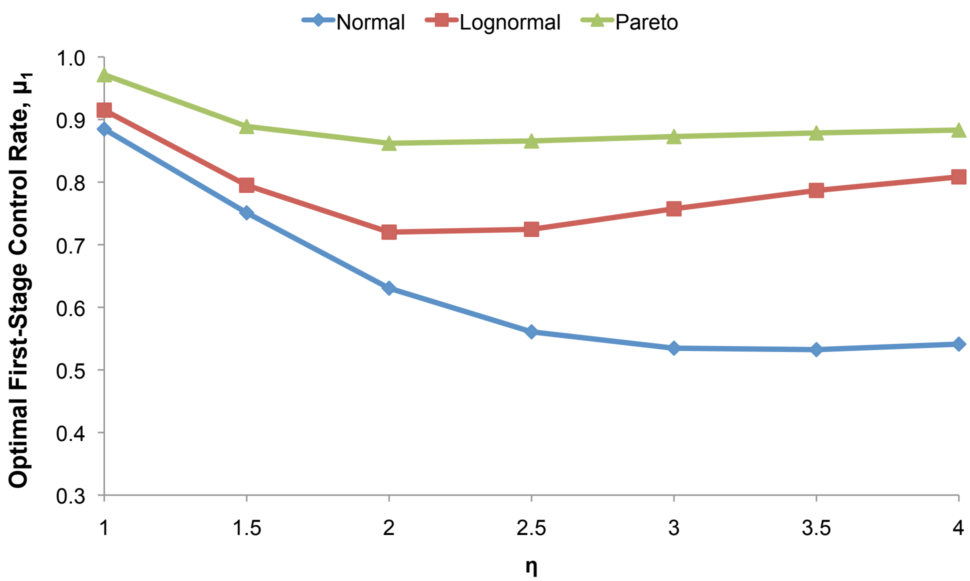

Figure 7.5: Optimal first-stage control rates (μ_1) as a function of the elasticity of marginal utility (η).

7.4 Discussion

The results from Section 7.3 suggest that answers to policy questions about vast uncertainty about uncertainty hinge critically on judgments about model representations of uncertainty, learning, damages, and discounting.[27] It is not merely uncertainty

[27]From a policy perspective, the results highlight the importance of flexibility, information gathering, and adjusting to new and updated information. Such flexibility provisions are central features

that influences climate policy decisions but also uncertainty about the evolution of uncertainty over time. Learning reflects the relationship between the arrival of new information and its effect on uncertainty, which can manifest itself in any change in the joint probability density function. Such information can come from a variety of sources like modeling, experimentation, and observation.

The results in Section 7.3 should refocus attention to the literature on the dynamics of learning, which offers many insights but unresolved questions about the nature and relevance of learning. A summary of learning in the context of global change and environmental policy can be found in Parson and Karwat (2011) and O'Neill et al. (2006). There is a more limited body of research specifically devoted to learning in relation to the climate sensitivity parameter, which is the central uncertainty in this chapter. Le Treut et al. (2007) trace the trajectory of climate sensitivity estimates over time. Many papers (Zaliapin and Ghil, 2010; Hannart, Dufresne, and Naveau, 2009; Stainforth et al., 2005) describe reasons for the currently wide range of climate sensitivity parameter estimates (e.g., cloud processes and oceanic response) and how increased research in these areas may reduce uncertainty. There are also many explanations for why the current distribution associated with the climate sensitivity is likely skewed and fat-tailed (Weitzman, 2012; Roe and Baker, 2011).

In problems of global change, there is a concern that the timescales associated with learning and adjustment may be long relative to the timescales over which policy-relevant decisions are made (Baker and Roe, 2009; Oppenheimer, O'Neill, and Webster, 2008; Allen and Frame, 2007). For climate change, this timescale incompatibility is related to "stock-accumulation inertias" in a range of "physical and biological processes that are extremely slow to respond to attempts at reversal" (Weitzman, 2012). For instance, deep ocean adjustments take several centuries to come to equilibrium; however, this timescale is weakly constrained given the limited understanding about ocean adjustment processes (Baker and Roe, 2009). These long inertial lags and pipeline commitments are the primary justifications in Weitzman (2012) for the 150-year expected learning time.[28]

of adaptive environmental management.

[28]Expectations about the role of corrective, recourse actions are shaped not only by the potential for learning but by barriers to reacting. Thus, the uncertainty associated with learning presents

The literature on the expected time to reduce uncertainty for the climate sensitivity parameter addresses these concerns predominantly by using Bayesian learning frameworks. This literature is extremely relevant to the results in this chapter, since Section 7.3.3 shows that optimal near-term abatement strongly depends on this expected learning time. Kelly and Kolstad (1999) show that it may take 90–160 years to learn the true value with 95 percent confidence and 110–200 years with 99 percent confidence. They describe the tradeoff between the benefits of accelerated information gathering and controlling emissions under uncertainty. Leach (2007) extends this work by showing that additional uncertainty in the persistence of temperature deviations may extend the required learning time by hundreds of years. Webster, Jakobovits, and Norton (2008) demonstrate how additional sources of uncertainty in climate processes can increase the time needed to reduce uncertainty and how observations of additional climate-related variables like sea-level rise can reduce learning times. The paper calculates that these additional observations may reduce uncertainty by 20–40 percent over the next 20–50 years. Kelly and Tan (2013) investigate the influence of fat-tailed uncertainty on the learning rate. They find an average learning time of about 80 years but show that probability mass in the tail of the distribution diminishes relatively rapidly if the true climate sensitivity value is not high, which means that it may be possible to learn enough to rule out very high values in the coming decades (even if the exact value is not known for much longer).

The important, decision-relevant characteristics of the learning process are how the dispersion of a probability distribution changes as new information arrives, the rate of new information arrival, and whether learning will converge to the true value. Two common assumptions about decision-making under uncertainty are that new research and information will lead to more accurate beliefs and that such information will reduce uncertainty. Although these intuitions hold in many decision contexts, it is not necessarily true *a priori* that learning will converge on the true value and/or reduce uncertainty in all cases (Oppenheimer, O'Neill, and Webster, 2008; Henrion and Fischhoff, 1986).

---

both scientific and political questions, as midcourse corrections require an understanding of climate system dynamics (i.e., to extract clear signals in time to act) and of political feasibility (i.e., to adjust emissions accordingly).

Disconcerting learning refers to the notion that new information can lead to larger uncertainty (Hannart, Ghil, and Dufresne, 2013).[29] In contrast, *reassuring learning* corresponds to the conventional idea that learning reduces uncertainty, which is assumed in many models of climate-related learning (Webster, Jakobovits, and Norton, 2008; Keller, Bolker, and Bradford, 2004). The potential for disconcerting learning is especially relevant in the context of the climate sensitivity parameter, since Hannart, Ghil, and Dufresne (2013) show that the two most prominent characteristics associated with a greater likelihood of disconcerting learning are high skewness and fat tails in the prior. A combination of these features, as with the Pareto distribution, results in a substantial increase in the likelihood of disconcerting learning and consequently in greater uncertainty about the future trajectory of uncertainty and the expected learning time.[30] For instance, Morgan and Keith (1995) asked climate scientists to estimate the probability that the uncertainty surrounding the climate sensitivity parameter would increase by 25 percent or more after a 15-year research program with an annual budget of $1 billion. The responses ranged from 0.08 to 0.30 (with an average of 0.19), as respondents viewed unforeseen complexities arising from research and experimentation as more likely to increase uncertainty. Although the policy implications of disconcerting learning in the context of climate change are unclear, distributions associated with disconcerting learning generate reassuring trajectories for most cases, and when such unexpected events arise, disconcerting learning is a transient state and will eventually result in reduced uncertainty (Hannart, Ghil, and Dufresne, 2013).[31]

Negative learning occurs when new information causes current beliefs about an

[29] Hannart, Ghil, and Dufresne (2013) provide a mathematical definition for the occurrence of disconcerting learning. The paper uses standard deviation as a measure of uncertainty, though other metrics are possible.

[30] For instance, uncertainty may increase in the near future as computing power permits large-ensemble simulations of increasingly complex coupled atmosphere-ocean general circulation models (Rowlands et al., 2012).

[31] Disconcerting learning reinforces the notion that constant or increasing uncertainty can be compatible with scientific knowledge accumulation (e.g., an improved understanding of the climate system). Consequently, although there are many unresolved questions about definitions of learning and scientific progress (Kitcher, 1993; Kuhn, 1962), uncertainty is not an appropriate metric to assess such advances.

unknown parameter to diverge from its *a posteriori* true value (Oppenheimer, O'Neill, and Webster, 2008).[32] *Progressive learning* refers to the increasing correspondence between an estimate and the true value over time. Apart from a few notable exceptions (Oppenheimer, O'Neill, and Webster, 2008; Keller and McInerney, 2008; Small and Fischbeck, 1999), most Bayesian analyses assume that the underlying model, including the likelihood function, are correct. However, the omission of a structural feature in a model can lead attempts to learn about an uncertain parameter to narrow on an incorrect value. Oppenheimer, O'Neill, and Webster (2008) show how such manifestations of negative learning can lead to substantial losses if, for instance, a climate model is subject to structural error and neglects a positive feedback on radiative forcing. Historical examples of negative learning include stratospheric ozone depletion, melting of the West Antarctic Ice Sheet, and energy and population projections (Oppenheimer, O'Neill, and Webster, 2008). Like disconcerting learning, negative learning has the potential to lengthen the duration of the learning path and to delay midcourse corrections.

Although this taxonomy of learning does not provide unambiguous implications for policy, a few modeling recommendations for the treatment of uncertainty and learning emerge from this literature. First, uncertainty quantification should be given increased research attention commensurate with its importance in determining model results. The overconfidence effect is a contributing factor in both negative and disconcerting learning, which suggests that more careful attention must be given to assuaging cognitive biases and to employing probability elicitation techniques that are not based solely on historical values, which may be unreliable for prediction. Greater focus should likewise be placed on poorly understood scenarios in the extremes of distributions and on alternate models, especially when such efforts generate prospective information with the possibility of changing the default decision. Oppenheimer, O'Neill, and Webster (2008) suggest that characterizing uncertainty should emerge as a "co-equal partner with consensus building." Second, given past experience and the possibilities for negative and/or disconcerting learning, modelers should compare

[32]There are many possible manifestations of negative learning, including increasing variance when a distribution is centered on the true outcome, decreasing variance centered on a value sufficiently far away from the true value, or shifting a distribution away from the true value.

forecasts with evolving observations to determine trends in estimation errors. Tools from disciplines like statistics, decision analysis, and the history of science can be used to understand how scientific judgments are made and to ascertain why errant forecasts were wrong (Craig, Gadgil, and Koomey, 2002).[33] Understanding the dynamics of error is important in understanding the dynamics of learning. Learning research could offer spillovers to other areas of inquiry owing to the prevalence of decision problems in which provisional, near-term commitments must be made but revised over time in response to new information.

This brief survey of the learning literature suggests that the simplified representation of learning described in Section 7.2 is almost certainly wrong. However, the existing literature does not unambiguously point toward a correct way to conceptualize or model learning in this context. The trajectory of uncertainty for values like the climate sensitivity parameter is uncertain. In highlighting the role of learning as a paramount feature in the debate about uncertainty, the analysis from this paper provides more questions than answers and replaces one intractable problem with another. "The expansion of our knowledge has expanded the circumference of our ignorance" (Fairbank, 2006). However, this analysis has offered critical insights for decision-making and identifies important dynamics that should be given greater research attention, which are important contributions in a research enterprise as complex as climate change economics.

7.5 Summary and Extensions

This analysis demonstrates how optimal policy prescriptions from IAMs are highly sensitive to the specification of uncertainty and learning. Given that the representations of these elements are key determinants of modeling results, analyses that do not test over a range of assumptions about the characterization of uncertainty and

[33] These dynamics are especially significant given that scientific estimates in global change research tend to be biased toward caution. Brysse et al. (2013) suggest that dynamics of the scientific method and assessment processes create systematic biases toward underestimation due to the methodological standards, epistemic procedures, and norms (e.g., objectivity, skepticism, restraint, moderation) of the scientific community.

its evolution may be misleading.[34] Weitzman (2009) has been influential in advancing the dialogue surrounding the role of uncertainty in energy modeling; however, a more realistic refinement of this work gives insights into the relative importance of uncertainty and learning under various specifications of discounting and damages.

Model results illustrate the value of learning and midcourse corrections on reducing consumption risks imposed by uncertain damages from climate change. The results demonstrate the value of early information about the severity of climate change, which echoes conclusions of previous dissertation chapters. It is not necessarily learning about particular outcomes that is important but the expectation of obtaining more information, which allows decision-makers to adjust their decision strategies. The potential for learning reduces the stringency of precautionary mitigation even under fat-tailed uncertainty. If perfect information about the climate sensitivity parameter were available immediately, it would reduce climate insurance through abatement by nearly 50 percent, and the expectation of learning in 80 years reduces control rate by 26–36 percent.[35]

Ultimately, fat tails impact near-term policy significantly when damages are strongly convex and when learning is slow (specifically, when information about the true climate sensitivity parameter does not arrive until after 2100). Thus, fat tails *per se* are not sufficient enough to merit immediate and stringent mitigation, as they also require reactive damages and slow learning.[36] The sensitivity of IAM results to the specification of the damage function hints at the need for more research on the economic impacts of large warming. The *terra incognita* of evaluating welfare implications in a

[34]Currently, testing the robustness of results over assumptions about discounting is a common practice and a prerequisite of credible analysis. Uncertainty and learning should be treated in a similar manner.

[35]The level of actual abatement embodied in future policies will likely be determined by political decisions tempered by many factors not accounted for in this simplified modeling framework. However, to extent to which scientific evidence, economic and ethical reasoning, and decision theory are used to inform decisions, climate policy should be guided by long-term risk management considerations, which include the possibilities of fat-tailed uncertainty and learning.

[36]The results demonstrate how incorporating fat tails does not suggest that we are on course for an inevitable (or even likely) disaster unless dramatic and immediate abatement is undertaken. Most catastrophic events are extremely unlikely to materialize, since tail probabilities are still comparatively small, even for the fat-tailed Pareto distribution. However, the analysis also suggests that the possibility of fat-tailed uncertainty justifies concern about avoiding low-probability extreme events by purchasing some degree of insurance against risks of severe outcomes.

world with temperature changes of 10 °C (28 °F) implies that, instead of using scarce resources to substantiate consensus issues related to central tendencies, the scientific and economic communities would be better served by investigating the boundaries of our understanding about extreme outcomes.

The model used for this analysis is a grossly simplified representation of many interconnected, complex, and uncertain systems. The analysis provides a stress test of robustness through a simple, transparent framework. The next step is to add extra complexity through the possibilities of partial and continuous learning, endogenous savings rates, more decision stages, and greater disaggregation to verify the conclusions from this simplified model. The broad qualitative insights about the sensitivity of policy recommendations to the specifications of uncertainty and learning, *mutatis mutandis*, would likely be similar in a model with these enhanced features, though the magnitudes of the interactions would be different.

An important research need is to investigate the dynamics of scientific learning. As described in Section 7.4, historical examples in global change research suggest that misjudgments and negative learning may have provided misleading advice to decision-makers during policy debates (Crutzen and Oppenheimer, 2008; Oppenheimer, O'Neill, and Webster, 2008). Future work should investigate how new information and the ostensibly best-available evidence may prove to be in error and how errant beliefs may actually influence policy-related decisions (Oreskes and Conway, 2010; Crutzen and Oppenheimer, 2008).

Viewed through the lens of energy modeling, the debate surrounding Weitzman (2009) implicitly centers on higher-order decision procedures for determining a suitable modeling framework in which to evaluate alternative climate policies given the possibility of fat-tailed uncertainty. The research in this chapter introduces a useful metric, the value of the stochastic solution, to clarify the conditions under which the deterministic benefit-cost analysis framework is an inadequate and potentially misleading decision-support tool. Using a sequential decision-making model, the results demonstrate how the conceptual apparatus of benefit-cost analysis can accommodate fat-tailed uncertainty with a few computational modifications. Although the analysis assumes that these stochastic modifications adequately represent risk, it can be argued that the uncertainties associated with climate change require another type of modeling framework with a different optimization criterion, especially if potential discontinuities, irreversibilities, and tipping points are integrated into the analysis (Hall et al., 2012; Kunreuther et al., 2012; Morgan et al., 2009; Dessai and Hulme, 2004; Kann and Weyant, 2000). Future work should investigate a more systematic method for determining suitable decision procedures for adjudicating between policy alternatives in different decision contexts with varying degrees of uncertainty (and uncertainty about uncertainty).

Chapter 8

Conclusions and Future Research

8.1 Summary of Findings

This dissertation explores several applications of sequential decision-making for energy problems under uncertainty. An overarching conclusion is that uncertainty analysis should be given more careful consideration during model development, calibration, and communication and should not be viewed as an after-the-fact, *post-hoc* analysis tool (Morgan and Henrion, 1990; Dowlatabadi and Morgan, 1993). Methods and metrics similar to those employed here should play more prominent diagnostic roles in energy modeling and integrated assessment modeling. Metrics like the value of the stochastic solution (VSS) and expected value of perfect information (EVPI) can be useful guides for prioritizing research efforts and for placing values on reducing the most consequential uncertainties. These values can be approximated using standard deterministic models, which may prove to be a useful, comparably low-effort way of directing attention toward quantifying and analyzing uncertainties that materially influence near-term decisions. In this way, such methods can be used as complements to common uncertainty analysis techniques (e.g., sensitivity and scenario analysis) as catalysts for determining the degree to which more uncertainty analysis is warranted in specific context.[1]

[1] Stochastic programming frameworks also can be complementary to robust decision-making analyses by suggesting candidate decision strategies.

A primary research objective is to investigate what insights can be learned through sequential decision-making approaches like stochastic programming, which would not be apparent using deterministic approaches. The range of applications in this dissertation suggests that the impacts of explicitly incorporating uncertainty may be large across a range of domains and that such methods may yield novel, decision-relevant results. For instance, the research and development (R&D) management work in Chapter 6 shows how wait-and-see valuation approaches, by not explicitly accounting for diffusion-related uncertainties like climate policy, may undervalue the hedging potential of technologies like nuclear.

Although stochastic programming frameworks can be computationally complex and data-intensive, these barriers may be small in comparison to potential gains, especially when spillovers related to uncertainty quantification and model diagnostics are taken into account. The usability and accessibility of these methods have increased significantly in recent years by a combination of advances in computing technology, improved analytical methods for optimizing mathematical programming models, and more readily available commercial solvers. This dissertation demonstrates the benefits of bridging state-of-the-art operations research with energy models to provide prescriptive guidance in making deliberative, informed choices. These analytical and computational advances can be leveraged to take into consideration a wider range of potential futures and to hedge against negative outcomes.

Another research goal is to link empirical assessments of uncertainty with modeling frameworks to provide normative decision support. Using a range of methods to quantify uncertainty (as detailed in Chapter 3.4), the results illustrate the material importance of input distributions and suggest that uncertainty quantification should be a co-equal partner with model building. For instance, the capacity planning results in Chapter 4 using outdated distributions illustrate the importance of model assumptions about a decision-maker's expectations and the significance of updating model data with the best-available information. Uncertainty quantification efforts may have important spillover benefits in other modeling domains, which increase the value of such endeavors.

Even when formal uncertainty analysis is performed, the results of this research suggest that probability distributions from existing studies often exhibit overconfidence and do not reflect a full range of possible outcomes. Overconfidence refers to the systematic tendency to underestimate uncertainty so that the subjective confidence of decision-makers in their probability assessments is higher than their objective accuracy (Fischhoff, Slovic, and Lichtenstein, 1977). This effect is seen in the dissertation in diverse areas like forecasts for natural gas prices (Chapter 3.4), predictions for the performance of gas-turbine-based units (Appendix B), and elicitations for energy technology costs (Appendix B). Since the overconfidence effect narrows probability distributions, metrics like the EVPI and VSS are consequently biased downward (Hammitt and Shlyakhter, 1999). This pervasive bias suggests that existing analyses likely underestimate the value of gathering information about unknown quantities and of explicitly accounting for uncertainty in modeling efforts.

Given this context of overconfidence, decision-makers should expect (at least in the near term) to observe an increase in uncertainty over time, suggesting an imperative toward epistemological humility and embracing the ostensibly paradoxical stance of expecting to be surprised. This point underscores how decision-making in many energy and climate domains involves an irreducible component of near-term uncertainty (Maslin and Austin, 2012). Decision-makers should not anticipate that perfect information will arrive and eliminate uncertainty in the near future. For problems replete with complexity, heterogeneity, and sparse data, a corollary is that decision-makers and policy-makers should resist the temptation to rush toward an unwarranted sense of certainty with the limited evidence supplied by a single publication. Such overgeneralizing and biased assimilation are distractions from more productive tasks associated with characterizing and analyzing actual decision problems, which take place against a background of persistent uncertainty. Decision-makers and policy-makers should update probability estimates and decisions in accordance with the measured pace of Bayesian inference and should use a principled, strategic approach to gathering and incorporating new information to iteratively revise beliefs. Decisions that rely on predictions related to, for instance, the climate sensitivity parameter, upstream emissions from unconventional gas production, and regional impacts from climate change should be attentive to these considerations.

In the longer term, energy and environmental modeling will benefit from adopting formal methods of uncertainty quantification sooner, as the use of the frameworks and metrics from this dissertation can reduce the deleterious effects of overconfidence and surprise. Future modeling efforts should carefully consider the impacts of potential surprises (even though such surprises, by their very nature, are elusive) and account for a broader range of uncertainties. The strategic selection of a wider array of sensitivities, robustness metrics, and probability distributions allows audiences to develop more complete insights from modeling exercises while also requiring more diagnostic experiments. Instead of allocating scarce resources to substantiate consensus issues related to central tendencies, the modeling community would be better served by investigating the boundaries of our understanding about extreme outcomes and surprises, as the analysis of climate policy under fat-tailed uncertainty in Chapter 7 illustrates. Mason (1969) highlights the importance of "counterplans" and disconfirming evidence for offering alternate model frameworks and assumptions in strategic planning settings. Given the complexity associated with energy and environmental systems, countermodels can provide maximal stress tests to existing models and can help to expose embedded assumptions. This dialectical scheme, Mason (1969) proposes, will assist decision-makers and modelers in synthesizing the jarring and incompatible viewpoints of different models and their proposed strategies into one that includes and transmutes them.

There is no one-size-fits-all strategy for analyzing uncertainty. The degree to which uncertainty is analyzed and incorporated in a given context depends on many factors like the risk exposure of decision-makers, resources available for analysis, and potential recourse actions. A range of quantitative and qualitative approaches should be considered and implemented as appropriate to meet the needs and practical requirements of decision-makers. Modelers should constantly be alert to the possibility that a simpler analysis approach is available to bound or approximate a quantity of interest, which can determine whether more sophisticated approaches are necessary.

Sequential decision-making should ideally be used in policy formulation and analysis. However, given the difficulties associated with implementing and formalizing such procedures in national and international policy, the greatest impact of these approaches may be through modeling to inform policy debates, to diagnose blind spots, to understand incentives of decision-makers, and to avoid bad decisions. A central objective of these approaches for uncertainty quantification and analysis is to promote engagement between decision-makers, policy-makers, and modelers using the best-available evidence. Stochastic programming efforts like those found in this dissertation are successful when they can promote discussions that would not otherwise happen. Many insights and relationships from more complicated frameworks likely could be represented in simpler models, but the effects could not be diagnosed and identified without building and using more complex models first.

8.2 Capacity Planning under Uncertainty Results

Chapters 3 through 5 investigate the dynamics of capacity planning and dispatch in the United States (US) electric power sector under a range of technological, economic, and policy-related uncertainties. Model results suggest that the two most critical risks in the near-term planning process are natural gas prices and the stringency of climate policy. As with other applications in the dissertation, this chapter suggests that planners are likely underestimating the impacts of uncertainty. The appreciable VSS values suggest that the widely used practice of approximating stochastic programming models by using *ad-hoc* combinations of deterministic model runs is suboptimal.

Stochastic strategies indicate that nuclear and lower-cost wind are strong candidates for near-term hedges against a variety of uncertainties while allowing the electric power sector to keep pace with growing demand and retirements in the coming decades. These technologies are attractive investments due to their low life-cycle greenhouse gas emissions (which reduces exposure to the climate policy uncertainty) and low fuel price volatility (relative to alternatives like natural gas), which means that these technologies are economical under a wide range of contingencies and are not as likely to be mothballed or decommissioned once new information becomes available. However, the high EVPI found in this analysis suggests that there is a limited availability of the hedging options in the electric power sector. Consequently, a takeaway for US policy-makers is that, if they can offer a higher degree of climate policy certainty, there will likely be appreciable economic benefits for the power sector.

In addition these hedging insights, the results demonstrate that delaying some investments and waiting for more information can be optimal. Model results suggest that the threat of stranded abatement investments outweighs precautionary effects and results in a propensity to delay near-term expenditures. In particular, the stochastic approach will avoid near-term investments in carbon capture and storage (CCS) due to the possibility that these assets would be decommissioned either if the climate policy too stringent or too lax or if public opposition prevents cost-effective carbon dioxide storage. These dampening effects of uncertainty are explained in terms of the optionality of investments in the power sector, leading to more general insights about uncertainty, learning, and irreversibility.

Consequently, utilities have little near-term incentive to build CCS-equipped capacity given uncertainty in climate policy. If learning effects are important for reducing costs to enhance CCS readiness in future decades, large-scale CCS deployment may require public-private partnerships for early pilot and demonstration projects as well as for R&D for capture systems with lower parasitic losses. A parallel analysis with limited CCS availability suggests that the value of CCS readiness in the second stage is $13.8 billion. This result illustrates that, although they are not ideally suited for short-term deployment, CCS technologies may be an important part of the long-term generation mix. More generally, the large losses under runs that exclude entire groups of technologies speak to the importance of maintaining a diverse portfolio of generation options that, if needed, are ready for deployment at reasonable cost and performance levels.

The largest losses occur when decision-makers' beliefs depart from the best-available information either by using outdated distributions for fuel prices or by adopting optimistic beliefs about the ability to postpone a comprehensive climate policy. These results of misestimation underscore the importance of using distributions that incorporate actual data (i.e., instead of stylized, *ad-hoc* distributions) and, more generally, of updating model data. Such conclusions are especially relevant given the limitations of existing approaches for uncertainty analysis in utility resource plans, as described in Chapter 3.4.

These results suggest that a sequential approach to climate policy (e.g., by implementing a new source performance standard) could incentivize preemptive and supererogatory abatement efforts until more comprehensive climate legislation is in place. These policies may be effective instruments to reduce cost risks for utilities, to safeguard against the erosion of public confidence in political institutions, to demonstrate the feasibility of emissions reductions, and to lower the probabilities of environmental hazards for society at large.

The modeling results offer many policy-relevant insights about the future role of unconventional natural gas in the US power sector. The value of control for upstream emissions from shale gas is shown to be substantial. Limiting methane leakage allows more gas units to be built and operate during the second stage in scenarios where higher carbon taxes are realized and natural gas prices are low to moderate. This research refocuses the debate about methane leakage onto making decisions under inevitable near-term uncertainty. This reconceptualization resists the notion that a single "true" value of leakage exists, since drilling sites and practices are heterogeneous and may change over time. The implication is that policy-makers should avoid unnecessary generalizations based on a single empirical study of a specific location and also should acknowledge that policy-making will always be shrouded in some degree of uncertainty. Even if the leakage value appears to be on the high end of the range, natural gas should not necessarily be removed from consideration as an abatement option, as control and capture technologies may become important elements in the abatement choice set.

Additionally, questions about whether a carbon price will increase or decrease natural gas consumption and whether shale gas availability will influence investments in renewable technologies are shown to hinge on interactions between gas prices and climate policy uncertainties. The shale gas boom will not impede long-term investments in low-carbon technologies if a sufficiently stringent climate policy is enacted in the coming decades. However, if policy-makers fail to provide suitable incentives for firms to internalize climate-related externalities, utilities may overinvest in gas-related infrastructure and underinvest in low-carbon technologies relative to their socially optimal levels. Such effects illustrate the importance of modeling the interactions between multiple uncertainties simultaneously.

8.3 Energy Technology R&D Portfolio Management Results

Chapter 6 presents a framework for allocating investments across a portfolio of energy technology R&D programs. One contribution of this research is to offer guidance about how R&D success valuations vary in different decision-making settings. Although theoretical results suggest that the effect of decision-making approaches on R&D valuation is equivocal, experiments using a two-stage stochastic model indicate that R&D is more valuable in second-best planning environments like when decision-makers use expected-value approaches. Additionally, these comparisons suggest that traditional, wait-and-see R&D valuation approaches likely overestimate the value of R&D success for many programs but undervalue the optionality and hedging potential of technologies like nuclear relative to the stochastic approach.

R&D plays an important insurance role in second-best policy environments. The results demonstrate how the expected value of R&D success is modulated by the degree to which environmental damages are internalized. When policy is less stringent than the socially optimal level, investments in R&D can mobilize low-carbon technologies that would not otherwise be cost-competitive.[2] In this way, R&D provides a secondary method of greenhouse insurance (Manne and Richels, 1993; Blanford, 2006; Baker and Solak, 2013) alongside mitigation for reducing exposure to risk, which is especially valuable if the timing and/or stringency of climate policies are suboptimal. For public R&D decisions, the results reinforce the importance of accounting for economic benefits accruing to all impacted stakeholders (and not just to innovating firms) and also stress the normative rationale of technology-push policies like R&D to complement demand-pull measures. Technology policy, even under appropriately

[2]R&D alone is an inefficient means for achieving emissions reductions. The environmental economics literature suggests that, if socially optimal abatement policies prove to be politically infeasible, costs induced by such barriers are likely to be substantial (Fischer and Newell, 2008).

stringent environmental controls, is valuable for reducing compliance costs, accelerating the arrival of commercially viable technologies, and preventing excess reliance on conventional technologies during the time before large-scale market penetration of more environmentally benign alternatives.

The proposed framework for R&D portfolio management offers a structure and set of tools for facilitating information gathering and model improvements. It is designed to offer support for difficult but necessary decisions in a way that makes its underlying dynamics and assumptions transparent but realistic. This framework can be updated and expanded as better information becomes available, as more uncertainties and R&D alternatives are identified, and as improved models are created. These traits are integral to the enterprise of managing technological change, as the characterization of innovation requires expert elicitations and other modeling judgments in areas where reasonable analysts and stakeholders may disagree.

Ultimately, a broad objective of this work is to assist decision-makers in thinking through the benefits of long-term, sustained R&D portfolio expenditures and to demonstrate the value of such programs. Given the high potential payoff from R&D but also the political constraints associated with public R&D allocations, even opponents of R&D projects should consider whether prudent and persistent R&D investments can be worthwhile insurance against the potential for another "wasteful, ill-directed response to likely future crises" (Cohen and Noll, 1991). Crisis-driven spending and the cyclical divestiture of program funding are not conducive to the production of the useful knowledge and specialized human capital, which require long-term R&D efforts.[3] Given the relatively inelastic supply of scientists and engineers in the short run (Goolsbee, 1998), brief spurts of R&D funding may exceed the absorptive capacity of the industry and eventually may undercut the potential for increasing the scale of R&D output (i.e., by increasing wages of a fixed labor supply), which makes a strong case for steady and predictable R&D funding.

[3] R&D also may assist with making alternative technologies and proposals available for additional development and deployment when a crisis does occur, which can help to make such expenditures less inefficient.

8.4 Fat-Tailed Uncertainty, Learning, and Climate Policy Results

Chapter 7 demonstrates how climate policy prescriptions from integrated assessment models are highly sensitive to the specification of uncertainty, learning, and damages from climate change. Given that the representations of these elements are key determinants of modeling results, analyses that do not test over a range of assumptions about the characterization of uncertainty and its evolution may be misleading.[4] The results support the observation that the aim of integrated assessment modeling is not to find a single optimal global policy to be followed in perpetuity but to avoid grossly suboptimal policies and to revise policies as new information becomes available. This reframing echoes the sentiments of Morgan (2011), which concludes that it is misleading to describe the problem as that of finding an "optimal global climate policy" instead of examining "widely robust strategies."

Model results illustrate the value of learning and midcourse corrections on reducing consumption risks imposed by uncertain climate damages. The potential for learning reduces the stringency of precautionary mitigation even under fat-tailed uncertainty. For instance, if perfect information about the climate sensitivity parameter were available immediately, it would reduce abatement by nearly 50 percent, and the expectation of learning in 80 years reduces control rate by 26–36 percent. Ultimately, fat tails impact near-term policy significantly when damages are strongly convex and when learning is slow (specifically, when information does not arrive until after 2100). Thus, fat tails *per se* are not sufficient enough to merit stringent mitigation immediately, which also requires reactive damages and slow learning.

Even when fat-tailed uncertainty is included, it is important to note that the model and results do not suggest that catastrophic outcomes are likely to materialize. Tail probabilities are still comparatively small for the fat-tailed Pareto distribution. However, the model does assume that there is a small, nonzero probability of severe

[4] The decision-relevant impacts of fat-tailed uncertainty comprises a considerably smaller fraction of the current literature, and this selection bias has the potential to be misleading to policy-makers.

impacts and that this probability increases in atmospheric greenhouse gas concentrations. The analysis suggests that the possibility of fat-tailed uncertainty justifies concern about avoiding low-probability extreme events by purchasing some degree of precautionary insurance against these risks until we can obtain more information about their nature, severity, and likelihood.

The model from Chapter 7 is a simplified representation of many interconnected, complex, and uncertain systems. The analysis provides a stress test of robustness through a simple, transparent framework. The next step is to add extra complexity through the possibilities of partial and continuous learning, endogenous savings rates, more decision stages, and greater disaggregation to verify the conclusions from this simplified model. Additionally, this research focuses only on uncertainty related to the climate sensitivity parameter but does not incorporate other sources like scientific uncertainty (i.e., imperfect and incomplete understanding of the climate system), socioeconomic uncertainty (i.e., imperfect understanding of climate change impacts, distributional impacts on populations, and adaptive capacities of these societies).

Although the analysis assumes that the sequential decision-making model modifications adequately represent risk, it can be argued that the uncertainties associated with climate change require another type of modeling framework with different optimization criterion, especially if potential discontinuities, irreversibilities, and tipping points are integrated into the analysis (Hall et al., 2012; Kunreuther et al., 2012; Morgan et al., 2009; Dessai and Hulme, 2004; Kann and Weyant, 2000). Unlike the capacity planning model in Chapter 3, the uncertainties associated with climate change and associated responses put decision-makers in a realm closer to ambiguity, where relevant parameters are difficult to conceptualize let alone to quantify. Given that uncertainty is typically represented in an expected-utility framework (von Neumann and Morgenstern, 1944), when are alternative decision-making approaches preferable? Millner et al. (2013) use expert elicitations to show how some climate scientists' beliefs about the climate sensitivity parameter cannot be captured with subjective probabilities. This result means that existing elicitations may understate uncertainty, making the analysis in Chapter 7 especially relevant, and suggests that future work should explore additional frameworks for eliciting and incorporating uncertainty about uncertainty (Heal and Millner, 2013).

The chapter highlights the critical role of metacognitive deliberation in casting a critical eye toward forecast-driven policy prescriptions. Such domains require the examination of historical forecasts and analyses of error and systematic bias to give a better perspective on what modelers are likely getting wrong in existing analysis (e.g., the ubiquity of overconfidence), even though it is believed that such models represent the best-available evidence. The importance of metacognitive reflection about forecasting biases is shown in the complex nature of the scientific learning process, which has not conformed to expectations about converging on the true value and/or reducing uncertainty across a range of settings (Henrion and Fischhoff, 1986; Oppenheimer, O'Neill, and Webster, 2008; Hannart, Ghil, and Dufresne, 2013). This effect is also illustrated in the optimistic R&D portfolio recommendations for nuclear, where the conclusions are tempered by the recognition of the systematic historical overestimation of nuclear power deployment and underestimation of its associated costs (Grübler, 2010; Hultman, Koomey, and Kammen, 2007; Koomey and Hultman, 2007; Cohen and Noll, 1991). Overall, understanding the historical context of forecasting and model building in this domain is critical in understanding the present limitations of quantitative analysis.

8.5 Expert Elicitation Results

A consistent theme throughout this dissertation is the need for more and better probability elicitations, which should be high research priorities in the near term. In particular, expert elicitations are critical for uncertainties surrounding technologies in the energy modeling community, given that such assumptions drive results and models typically have time horizons of many decades. Although progress has been made in recent years (e.g., conducting elicitations for many technologies, making data widely available, and combining elicitation insights across research groups), energy technology expert elicitations remain highly active areas of research. For the integrated assessment modeling community, expert elicitations are important methods of ascertaining more comprehensive understandings of uncertainty compared with model-driven approaches, since experts have intimate knowledge of the strengths and limitations of existing models.

The results across many applications in this dissertation demonstrate how model results are sensitive to input distributions. For instance, the one-sidedness of the optimal R&D portfolios in Chapter 6 comes from the use of empirical values from elicitations and strongly suggests that modelers should prioritize nuclear power elicitations with careful debiasing. The importance of elicitations reflects a broader concern about more purposeful uncertainty quantification for input parameters across the energy and integrated assessment modeling communities. When undertaking uncertainty analysis, the use of *ad-hoc* distributions may ignore important dynamics of the system under investigation, which may give misleading insights. Chapter 7 explores the non-robustness of policy recommendations to the representation of uncertainty about the climate sensitivity parameter (and damage function) and underscores the importance of assessing probabilities associated with such parameters with resources commensurate with their influence on model results.

For energy technology elicitations, Appendix B examines the factors that enhance the reliability of these probability assessments and discusses unresolved questions about best practices for elicitation protocols. The elicitation for gas turbine systems shows how face-to-face assessments are extremely useful in critically examining reported probability values, particularly for the tails of the distribution. Based on these observations and the prevalence of the overconfidence effect, future research should examine to what degree at-a-distance elicitations exhibit greater overconfidence compared with in-person protocols and how interactive digital tools can bridge this gap if it is sizable. Answers to these questions are especially relevant given the need for more frequent elicitations involving rapidly changing technologies like solar (Reichelstein and Yorston, 2013), where it is important to use techniques that can save time and money while not compromising quality.

Ultimately, one of the largest benefits of the elicitation process is that it gives modelers more opportunities to consult technical experts who have the greatest experience and familiarity with technologies. Elicitations have an important role in energy modeling even in a deterministic setting to help modelers and decision-makers to identify and avoid potential blind spots in the planning process. For instance, exogenous technological progress in deterministic models is typically informed by engineering cost estimates, which should rely on elicitations to assess expert opinion and to structure sensitivities.[5]

In some respects, expert elicitations are as important for future modeling efforts as they are for those in the present. Probabilistic assessments preserve information about current beliefs for use in the future, which means that formally capturing such beliefs is necessary for hindcasting exercises. Therefore, elicitations play an integral part in constructing information management systems, improving models for decision support, and combating hindsight bias (Fischhoff, 1982; Fischhoff and Beyth, 1975). These assessments are likewise necessary for evaluating the dynamics of learning (Hannart, Ghil, and Dufresne, 2013; Oppenheimer, O'Neill, and Webster, 2008) and for understanding why errant forecasts were wrong (Craig, Gadgil, and Koomey, 2002).

Tools like expert elicitations for uncertainty quantification, combined with the uncertainty analysis techniques described in this dissertation, will be increasingly important amid an environment of substantial uncertainty. The emergence of increasingly powerful tools for coping with risk comes at an opportune time. The coming decades will exhibit pervasive, multifaceted uncertainty and simultaneously will require increased capital investments. According to the International Energy Agency (IEA, 2012b), global investments in electricity generation and distribution will exceed $17 trillion over the next two decades. Investments in the US power sector must replace an aging fleet of coal (and later nuclear) generators while keeping pace with growing demand for energy services and contending with a diverse and uncertain spectrum of investment alternatives. Additionally, policy-makers have a critical role in this context and may be able to exert some influence over the decision landscape through climate and technology policies. Suboptimal decisions can impact the US economy

[5] Additionally, Parson (2003) suggests that technological assessments and elicitations may serve a broader function by altering the "technological feasibility on which they are reporting, by advancing present technical skill, solving problems, and identifying and removing barriers to the implementation of new processes and products."

and global competitiveness; cost ratepayers, investors, and taxpayers considerable sums; and have significant environmental impacts.

This pervasive risk also represents an important opportunity. When uncertainty is prevalent and the long-term consequences of near-term decisions are imperfectly understood, hedging strategies allow decision-makers to shape available options, to cope with the unknown, to learn from errors, and to exploit new information as conditions change. The near-term planning opportunities for public and private decision-makers are especially valuable in light of the available time before new capacity is needed. In the interim, decision-makers may have the ability to direct R&D toward technologies and observe whether emerging technologies will exhibit the technological readiness and cost-competitiveness needed for large-scale deployment. As this dissertation emphasizes, timing is critical for the energy sector at this juncture.

The book *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis* (Lempert, Popper, and Bankes, 2003) succinctly describes this simultaneous challenge and opportunity:

> The biggest paradox is that our greatest potential influence for shaping the future may often be precisely over those time scales where our gaze is most dim. . . Where the future is ill-defined, unpredictable, hardest to see, and pregnant with possibility, our actions may well have the largest effects in shaping it.

Ultimately, the goal for decision-makers is to find more efficient and flexible approaches for grappling with risks associated with the unexpected. The approaches, metrics, and results presented in this dissertation offer support and insights that assist decision-makers toward those ends in problems under uncertainty.

Appendix A

Vector Autoregressive Model for Fuel Prices

This appendix presents details of an approach for estimating correlated probability distributions for natural gas and coal prices, which is briefly introduced in Chapter 3.4. First, motivations for this analysis are presented in Section A.1. Next, Section A.2 discusses model selection and development considerations, followed by a description of the data and estimation process in Section A.3. This appendix concludes with a presentation and discussion of the results in Section A.4.

A.1 Introduction

Prices for energy resources are uncertain and fluctuate based on many complex factors. In the case of natural gas, uncertainty about future prices is also driven by recent discoveries and increased domestic production of shale gas (Moniz, Jacoby, and Meggs, 2010). Although abundant gas resources suggest expanded use in the electric power sector, uncertainty about the environmental impacts of production and long-run production costs make the extent of this growth unclear (Huntington, 2013; IEA, 2012a; DOE/EIA, 2011a; Coleman et al., 2011). Additionally, natural gas price uncertainty will be influenced by the unknown policy environment, public acceptance of hydraulic fracturing, and uncertainty surrounding life-cycle emissions for shale gas.

Quantifying and understanding the uncertainty associated with forecasts for future fuel prices is critical for decision-makers. Many public and private choices, both in energy industries and other sectors of the economy where energy resources are factors of production, rely on estimates of energy prices and their associated uncertainty to inform decisions.

Given the importance of uncertainty quantification in this context, this chapter formulates and applies a vector autoregressive (VAR) model to estimate probability distributions over natural gas and coal prices for electric power generators. Using historical data for delivered fuel prices from the 2011 Annual Energy Review (DOE/EIA, 2011b) published by the Energy Information Administration (EIA) and forecast data from the 2012 Annual Energy Outlook (DOE/EIA, 2012), the model uses a two-step process to estimate the trend and variability for fuel prices and then employs this VAR model to create density functions for annual price growth rates. This modeling approach is based on the techniques developed by Zdybel and Baker (2013).

A.2 Model Selection and Development

The VAR model represents the linear interdependencies among multiple time series. The value of each variable at time t depends on its own past values and the past values of other variables. The reduced p^{th} order VAR, which is denoted VAR(p), can be expressed as:

$$y_t = c + A_1 y_{t-1} + A_2 y_{t-2} + \ldots + A_p y_{t-p} + e_t \tag{A.1}$$

where y_t is a vector $(k \times 1)$ of variable values at time t, c is a vector $(k \times 1)$ of intercept constants, A_i is a matrix $(k \times k)$ of coefficients (for all $i = 1, \ldots, p$), and e_t is a vector $(k \times 1)$ of error terms.

The error terms are assumed to satisfy the VAR definitions:

1. All error terms have mean zero: $\mathbb{E}[e_t] = 0$

2. The $k \times k$ positive-semidefinite matrix Ω represents the contemporaneous covariance matrix of error terms: $\mathbb{E}[e_t e_t'] = \Omega$

3. No serial correlation: $\mathbb{E}[e_t e_{t-k}'] = 0, \ \forall k \neq 0$

The first differences of log prices in the model can be interpreted as the log of price growth in each period. The selection of log price differences guarantees that forecast prices will be positive when generating price paths. Additionally, log prices reduce possible effects of heteroscedasticity. To model in differences, one must first verify that each time series is difference stationary and account for cointegration if it exists (i.e., if time series variables share a common stochastic drift), which is done in Section A.3.

The trend specification uses EIA projections from the 2012 Annual Energy Outlook and historical data from the 2011 Annual Energy Review. The choice to use the EIA's forecast data is motivated primarily by the objective of incorporating the most up-to-date projections, since the Annual Energy Outlook forecasts embody the best-available data about factors driving the future price trend. A secondary motivation is that using Annual Energy Outlook data permits model updating as better information becomes available. As shown in later sections, the resulting trend does not try to replicate the EIA forecast but instead incorporates it into the data.

The historical variance is applied to the future portion of the trend to estimate probability distributions for fuel prices. The model uses a decomposition method similar to seasonal decomposition (but uses past variability to define uncertainty).

The resulting model can be used as a decision support tool in many settings. First, the VAR model can be applied directly in a Monte Carlo framework to generate sample paths, which can provide decision-makers with a sense of risk over outputs of interest in an uncertainty propagation analysis. Second, a similar Monte Carlo analysis can be employed to generate probability distributions over random variables, which are subsequently used as inputs in a sequential decision-making framework. This application is used here, as described in Section A.4. Finally, the sample price paths can be used graphically as an aid for decision-makers to visualize the potential range of future uncertainty.

A.3 Data and Methods

A.3.1 Data

Historical data come from the EIA's 2011 Annual Energy Review (Tables 6.8 and 7.9). Prices between 1980 and 2011 represent the average delivered values for natural gas and coal in the electric power sector. Forecast data are from the EIA's 2012 Annual Energy Outlook and give prices for steam coal and natural gas between 2012 and 2035. Figure A.1 illustrates these log prices in first differences.

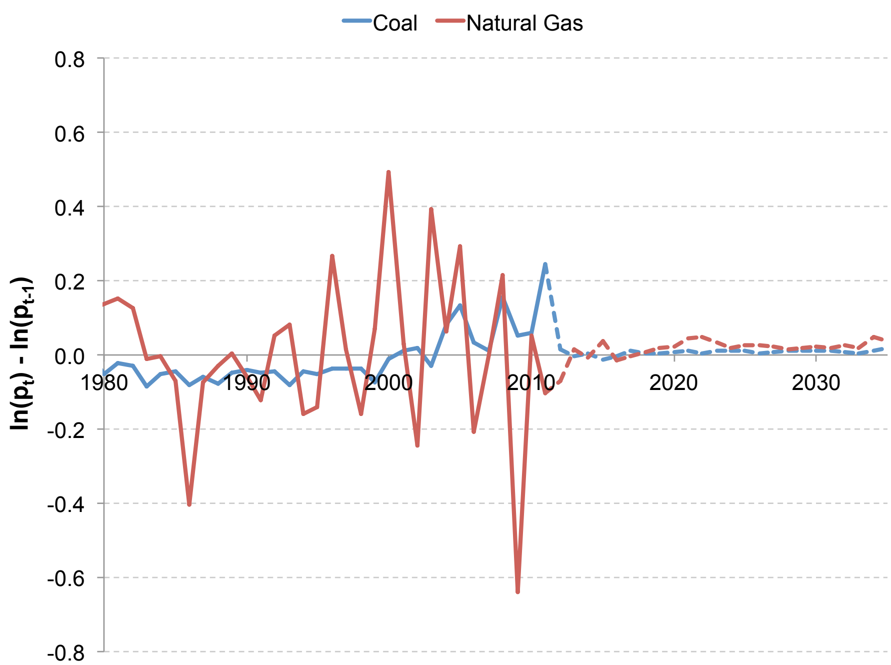

Figure A.1: Log prices in first differences for coal and natural gas. Historical fuel prices (left) are indicated by solid lines and forecast values (right) by dashed lines.

To assess the appropriateness of differencing the time series data, I first determine that prices in levels are nonstationary (since it is not appropriate to difference stationary data) and that they are integrated of order one. I use the augmented Dickey-Fuller test to check for a unit root in the time series sample. The null of a unit root is not rejected using log prices in levels but is rejected in first differences. This result suggests that prices are integrated of order one and that a cointegration test must be performed.[1] The Engle-Granger procedure test statistic indicates no cointegrating relationship, which suggests that a standard VAR model in differences is suitable for this application.

A.3.2 Determining the Number of Lags

The number of lags (p) for the p^{th} order VAR is determined using the Akaike information criterion (AIC).[2] As discussed in Enders (2009), the multivariate generalization of the AIC is:

$$AIC(p) = 2s + T \ln |\hat{\Omega}| \tag{A.2}$$

where T is the number of useful observations and $|\hat{\Omega}|$ is the determinant of covariance matrix of residuals for the estimated system.[3]

The AIC calculations suggest that a model without lags is the best choice for modeling the fuel prices considered here. Although not technically a VAR model, the model without lags can be viewed as a random noise process with contemporaneously correlated disturbances.

A.3.3 Two-Step Estimation Process

The first step is to determine the trend by estimating the VAR(p) model in Equation A.1 using historical and forecast values. This step gives values of c, A_i, and Ω (i.e., the covariance matrix of error terms). Here, c and A_i define the trend, and Ω represents uncertainty. The second step is to define the variability by re-estimating Ω based on the uncertainty of just the historical data in relation to the trend from the previous step. This variability calculation reflects only historical behavior by removing the effect of the forecast period.

[1] If a cointegration relationship exists, an error correction model must be used.

[2] Given the small sample size, a correction for finite sample sizes is used: $AIC_c = AIC + \frac{2s(s+1)}{n-s-1}$, where AIC denotes the AIC without the correction, s is the number of model parameters $s = k(1+pk)$, and n denotes the sample size.

[3] Residuals follow a multivariate normal distribution: $e_t \sim \mathcal{N}_m(0, \Omega)$.

A.4 Results

Using the two-step estimation process from the previous section, the VAR model with estimated parameters is:

- Coal (with standard error 0.0077)

$$\Delta \ln(y_t^c) = -0.0017 + e_t^c$$

- Natural gas (with standard error 0.0219)

$$\Delta \ln(y_t^g) = 0.0084 + e_t^g$$

- Covariance matrix

$$\Omega = \begin{vmatrix} 0.0033 & 0.0006 \\ 0.0006 & 0.0270 \end{vmatrix}$$

Since the covariance between coal and natural gas prices is small, these two random variables can likely be treated independently without a loss of fidelity, which is assumed in Chapter 3.4.

The historically adjusted covariance matrix is:

$$\Omega = \begin{vmatrix} 0.0038 & 0.0018 \\ 0.0018 & 0.0469 \end{vmatrix}$$

Thus, the variances increase when considering historical data only.

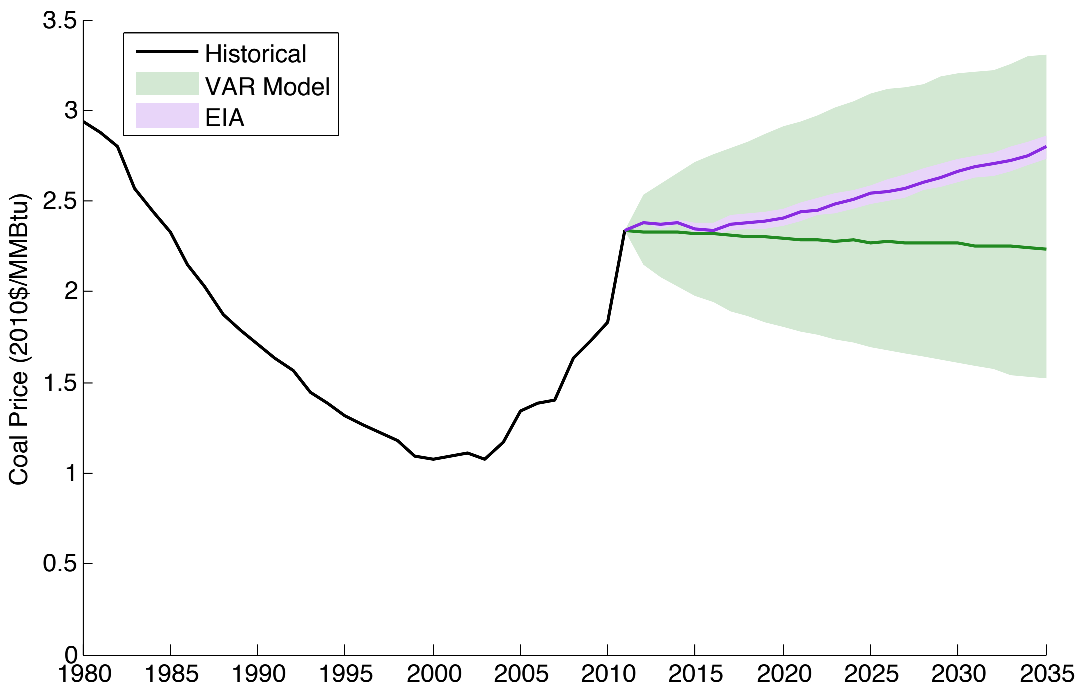

Figure A.2: Historical and forecast prices of coal for the electric power sector. The VAR model results show the 10^{th}, 50^{th}, and 90^{th} percentile values. The EIA cases represent the low price, reference, and high price scenarios.

Using the model with estimated parameters in a Monte Carlo simulation with 5,000 samples, Figures A.2 and A.3 show the uncertainty ranges for coal and natural gas prices, respectively. These figures show the 10^{th} and 90^{th} percentile bands for the correlated distributions along with the median values over time. The ranges from the price-path outcomes reflect fuel price uncertainty from a variety of sources like the economy, technological advances, policy, demand curve shifts, and other factors.

The trend for gas prices closely mirrors the EIA forecast, as the VAR model suggests that prices will increase slightly over the next couple decades. In contrast, the model results suggest that the uncertain range of prices may be much wider than the EIA currently projects, both on the lower and higher ends of the distribution.[4] Although the percentile ranges are broad for both fuels, these percentile bands are reasonable reflections of actual historical price behavior and are especially sensible

[4]These qualitative insights are similar for the coal analysis.

for probabilistic frameworks that consider low-probability, high-impact surprises.

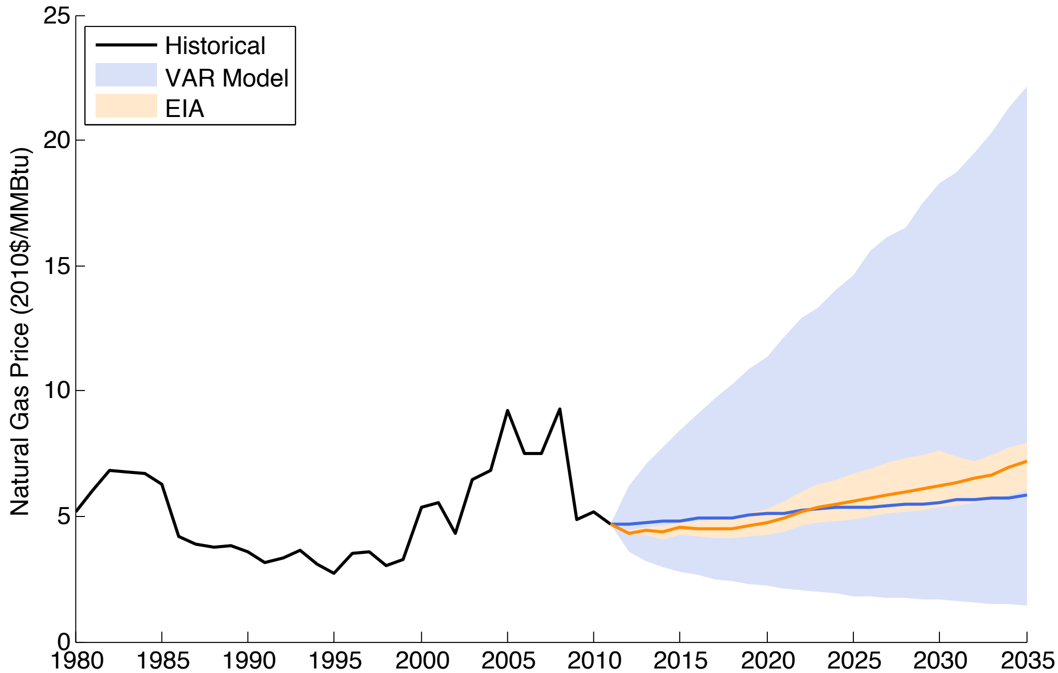

Figure A.3: Historical and forecast delivered prices of natural gas for the electric power sector. The VAR model results show the 10^{th}, 50^{th}, and 90^{th} percentile values. The EIA cases represent the low price, reference, and high price scenarios.

Modelers do not often quantify distributions over critical outputs or attach probabilities to possible scenarios, and there is evidence that, when analysts do quantify uncertainty, they tend to underestimate the range and probabilities associated with non-expected-value outcomes (Shlyakhtera et al., 1994). Although the EIA scenarios do not have associated probabilities, these results seem to support this observation and are consistent with the overconfidence effect (i.e., the cognitive bias where confidence intervals are assessed too narrowly) at an institutional level.

Appendix B

Energy Technology Expert Elicitations: An Application to Natural Gas Turbine Efficiencies

Expert elicitations play important roles in quantifying uncertainty about future cost and performance characteristics of energy technologies, as these estimates inform a range of decision and modeling efforts within the energy community. This chapter examines the factors that enhance the reliability of these probability assessments and discusses unresolved questions about best practices for elicitation protocols. These insights are applied in a case study to understand the current state of knowledge regarding the future of gas turbine systems for electricity generation. Elicitation results are used as inputs to the capacity planning model in Chapter 3.

The results support the conclusion that prospective efficiency increases are likely to be smaller than historical trends, which demonstrates the utility of elicitations in capturing dimensions of technical change that may be absent from forecasting methods that rely primarily on historical data. However, these median values are still appreciably higher than the efficiencies used in many integrated assessment models.

B.1 Introduction

Uncertainty analysis has played an increasingly prominent role in energy modeling in recent years (Kann and Weyant, 2000), particularly in regard to technological change. This focus comes as no surprise given that assumptions about how technologies evolve over time are leading determinants of modeling results (Weyant, 2004). Despite considerable unknowns about the dynamics of technological change, it is necessary to quantify this uncertainty about cost and performance metrics in a range of modeling settings. Obtaining a set of potential outcomes and some idea of their relative likelihoods is required no matter if uncertainty analysis is conducted implicitly (e.g., using sensitivity analysis or propagating uncertainty through deterministic models) or explicitly (e.g., through sequential decision-making frameworks like stochastic programming). The interest in characterizing technological uncertainty has grown in the presence of proposed energy and climate policies to manage technical change through research and development (R&D).

Although there are many formal methods of quantifying uncertainty, expert elicitations are uniquely suited for characterizing technological uncertainty. Statistical approaches that rely primarily on historical data may not contain sufficient information to form conjectures about the future progress or returns on research investments for specific technologies. Since breakthroughs are fundamentally unique, planners often cannot extrapolate past trends into the future or use relative historical frequencies to generate probability distributions over successes of technologies. Thus, when past data are unavailable or of limited use, one of the only remaining options is to ask individuals with expertise for their best professional judgments, which often take the form of expert elicitations (Morgan and Henrion, 1990).

An expert elicitation is a structured, formal process for collecting and assessing probabilistic estimates about uncertain quantities (Keeney and von Winterfeldt, 1989). These elicitations allow expert knowledge about specific technologies to be embedded in models instead of relying on stylized, *ad-hoc* distributions over parameters of interest, which may be selected with limited consultation about the current state of knowledge in a technological domain.

The objectives of this chapter are to:

- Survey the existing literature on energy technology expert elicitations
- Highlight the best practices and unresolved questions relating to existing elicitation approaches while suggesting fruitful avenues for future work
- Demonstrate effective elicitation techniques for an overlooked technology that merits greater attention—namely, natural gas turbine architectures for stationary power generation

B.2 Energy Technology Expert Elicitations

Considerable uncertainty about future states of energy technologies suggests that it is important to collect expert judgments about a range of possible outcomes instead of focusing only on central tendencies. In this setting, analysts cannot reliably assume that statistical analyses of historical trends or technological analogues (Rai, Victor, and Thurber, 2010) will provide accurate forecasts for the future evolution of energy technologies.[1] However, despite considerable uncertainty, probabilistic estimates from a diverse set of experts, encoded through a structured elicitation process, can offer valuable insights into technological developments.

B.2.1 Existing Work

Elicitations have been used for decades to encode the knowledge, judgment, and experience of experts in fields where uncertainty and risk are critical components of decision-making (O'Hagan et al., 2006; Morgan and Henrion, 1990; Staël von Holstein and Matheson, 1979). Since work by Tversky and Kahneman (1982), protocols for elicitations have been carefully designed using insights from psychology, decision analysis, risk analysis, economics, and statistics to reduce distortions from cognitive

[1] Frequently employed methods for projecting unit cost or performance characteristics using historical trends include regression analysis (Söderholm and Sundqvist, 2007), decomposition (Nemet, 2006), and monitoring for precursors (Martino, 1987).

biases and heuristics. Many researchers have investigated the strengths and shortcomings of various elicitation methods, and comprehensive overviews of the literature on the psychology of probability assessment and on elicitation approaches have been published (Kuhnert, Martin, and Griffiths, 2010; O'Hagan et al., 2006; Gathwaite, Kadane, and A., 2005; Meyer and Booker, 2001; Hoffman et al., 1995; Morgan and Henrion, 1990; Hogarth, 1975).

This chapter focuses on elicitation methods and applications for quantifying future cost and performance characteristics of energy technologies. The emphasis reflects the objectives of surveying current practices and unresolved questions in this policy-relevant area and also of applying these insights to investigate the future performance of gas-turbine-based technologies in the power sector. Although elicitations have been applied across a range of industries and research domains (Kuhnert, Martin, and Griffiths, 2010; Ayyub, 2001; O'Hagan, 1998; Hora and von Winterfeldt, 1997; Morgan and Keith, 1995), the application of elicitations to energy technologies began in earnest only recently. The limited research attention may come as a surprise given the pervasiveness of uncertainty in this domain and early interest in such analysis.[2]

For energy modeling, existing research uses elicitations to explore the future of several specific supply- and demand-side technologies. The most common objective is to inform questions of energy R&D policy, which has tremendous uncertainty about *ex-ante* returns on investments. The product of these elicitations is a rich set of data that encodes experts' best probabilistic judgments about future cost and performance characteristics for specific technologies conditioned on R&D effort and outcomes.

Table B.1 shows a non-exhaustive list of major energy technology elicitations in recent years. The five institutions conducting widespread elicitation research are across multiple energy technologies Carnegie Mellon University, the United States (US) Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE), Fondazione Eni Enrico Mattei (FEEM), Harvard University, and the University of Massachusetts Amherst.

- *Carnegie Mellon University*: Elicitations were largely conducted by researchers

[2]The Rasmussen report (NRC, 2010) on nuclear reactor safety is a prominent early example and the first to use quantitative expert judgments in a large risk analysis (Cooke, 2013).

affiliated with the Department of Engineering and Public Policy in a range of decentralized studies for amine-based carbon capture and storage (CCS) technologies (Rao et al., 2006), photovoltaic solar (Curtright, Morgan, and Keith, 2008), and small modular reactors (Abdulla, Azevedo, and Morgan, 2013).

- *Office of Efficiency and Renewable Energy (EERE)*: Researchers conducted elicitations for 40 renewable energy and efficiency technologies to support R&D portfolio management decisions using the Stochastic Energy Deployment System (SEDS) model, which has a Monte Carlo simulation framework. Affiliated researchers include Sam Baldwin (EERE), Max Henrion (Lumina), Thomas Jenkin (NREL), and Jim McVeigh (NREL).

- *Fondazione Eni Enrico Mattei*: Valentina Bosetti and colleagues have conducted elicitations for many energy technologies within a European context as part of the ICARUS project with a focus on the impacts of R&D (Bosetti et al., 2012; Fiorese et al., 2013).

- *Harvard University*: Laura Diaz Anadon and colleagues from the Energy Technology Innovation Policy Research Group within the Belfer Center for Science and International Affairs at Harvard's Kennedy School conducted elicitations in support of the research and publication of their *Transforming US Energy Innovation* report (Anadon et al., 2011).

- *University of Massachusetts Amherst*: Erin Baker and colleagues conducted elicitations for a variety of energy technologies, including nuclear (Baker, Chon, and Keisler, 2008), CCS (Baker, Chon, and Keisler, 2009b), solar (Baker, Chon, and Keisler, 2009a), battery technologies for vehicles (Baker, Chon, and Keisler, 2010), cellulosic biofuels (Baker and Keisler, 2011), and CCS energy penalties (Jenni, Baker, and Nemet, 2013).

Table B.1: Existing literature on energy technology expert elicitations.

	Carnegie Mellon	EERE	FEEM	Harvard	UMass Amherst
Supply-Side Technologies					
Nuclear			30	25	4
Coal with CCS	10			13	4
Gas with CCS				13	
Bioenergy and Biofuels			15	8	6
Solar	18		16	11	3
Wind					
Grid-Scale Storage				25	
Demand-Side Technologies					
Vehicles				9	7
Energy Efficiency				9	
Policy and/or R&D Scenarios	Yes	Yes	Yes	Yes	Yes
Elicited Years	2015 (CCS) 2030, 2050 (solar)	2015, 2020, 2025	2010, 2030	2010, 2030	2020, 2050
Year(s) Conducted/Published	2006–2012	2008–2010	2011–2012	2011	2008–2012
Protocol Method	Mail (CCS); combined mail/online, and face-to-face (solar)	Unknown	Combined online and group (nuclear); face-to-face (biofuels, solar)	Mail	Mail/online (nuclear, solar, vehicles), combined mail/online, face-to-face, and phone (CCS, biofuels)
Context	US	US	EU	US	US
Associated Model(s)	N/A	SEDS	WITCH	MARKAL	MiniCAM/GCAM

NOTE: Colored cells indicate, for a given research group, whether elicitations for a particular technology were not conducted (white), conducted (light orange), or conducted with published data (light blue). The values inside of technology cells indicate the number of experts included in the study (where available).

There have also been efforts to make elicitation results more accessible and to compare and aggregate their insights. Megajoule.org is a website spearheaded by Max Henrion for sharing and reviewing elicitation results. The Technology Elicitations and Modeling Project (TEAM) is developing an integrated framework for analyzing and communicating the results from energy technology elicitation efforts. A related collaboration between Harvard and FEEM researchers compares US and European Union (EU) elicitations for the future of nuclear power (Anadon et al., 2012).

B.2.2 Discussion of Unresolved Questions

Given the costly and time-consuming nature of elicitations, it is important to identify and understand the factors that enhance their quality and usefulness. This section highlights unresolved questions from the literature on energy technology expert elicitations and discusses the implications of these issues for modeling results based on such elicitations.

In-Person Elicitations, Conditioning, and Tail Events

Perhaps the most significant discrepancy between elicitation protocols is the method of administering elicitations and whether it is preferable to conduct them in person or at a distance.[3] Although at-a-distance methods are more economical and may allow greater participation, face-to-face elicitations have typically been preferred in the broader elicitation community due to a belief that such protocols yield higher-quality outputs (Curtright, Morgan, and Keith, 2008; Phillips, 1999).

One of the largest concerns about at-a-distance elicitations is that experts may be conditioning their responses on unspecified events. For instance, results of a recent elicitation for nuclear technologies (Anadon et al., 2012) demonstrates how experts believe that capital costs for Generation III/III+ reactors would be higher than at present. However, questions remain about whether this increase is due to forgetting curve effects, commodity price escalations, regulatory costs, or another random variable. Aggregate elicited values like price changes are causally overdetermined. It is impossible to decompose an expert's response to determine their beliefs about which factors influenced their response most without the ability to ask follow-up questions to determine what is implicitly being conditioned upon (e.g., depreciation of knowledge capital, increasing steel prices, inflation). Experts' mental models play central roles in the elicitation process (Morgan et al., 2002), but such models are inaccessible without the "interactive and iterative" (Jenni, Baker, and Nemet, 2013) feedback between the elicitor and the expert. Although feedback steps can mitigate some of these challenges for mail or digital elicitations, it is considerably easier to request feedback in an in-person setting and to reassess values immediately if it is discovered that the expert is conditioning on something that the interviewer does not intend.

One method of avoiding these omitted variable biases while retaining the convenience and cost reductions of at-a-distance elicitations is to make more use of innovative electronic techniques for conducting elicitations. Web-based interactive

[3] Although there is disagreement about how to conduct individual elicitations, there is broad agreement among energy technology research groups that individual elicitations are preferable to group methods. This sentiment aligns with recommendations in the elicitation literature, which caution against biases associated with group dynamics that can inhibit dissenting options (Meyer and Booker, 2001).

interfaces for authoring and hosting elicitations like Near Zero allow for more feedback from an expert conditional on their responses. In general, experimenting with newer elicitation techniques, particularly in ways that utilize digital tools and combine well-documented best practices from different methods, can improve the quality of elicitations over time. For instance, Anadon et al. (2012) use a novel, two-phase approach for conducting nuclear elicitations that begins with interactive online elicitations and a group meeting afterward.

A related issue surrounds the most effective means of assessing non-central probability estimates like the 10^{th} and 90^{th} percentiles. In the domain of energy technologies, the probabilities of extreme left-tail events (e.g., low capital costs resulting from technological breakthroughs, which may lead to wide deployment of a particular technology) and right-tail events (e.g., unexpectedly large costs that result from an inability to surmount engineering hurdles) are important to assess properly. However, assessing extreme values can be problematic owing to a host of cognitive biases, which impede careful consideration of low-probability events (Tversky and Kahneman, 1982). The most common bias is the overconfidence effect, which breeds underestimation of tail events. A failure to identify or correct overconfidence can result from not having an interviewer interact with and question an expert in real time, giving feedback about egregiously narrow distributions. Debiasing is particularly challenging for the overconfidence effect. Probability estimates may still exhibit this bias even when assessors are knowledgeable about its existence, which means that simply providing an information packet before elicitations may not be enough to safeguard against excessively narrow distributions.

Selection of Experts

The identification and selection of experts may be nearly as important as the design of the protocol itself. Although many technological elicitations are conducted to gain probabilistic information about future costs and performance characteristics, requesting cost and performance values from the same experts can be problematic. The catch-22 of technological elicitations is that experts must be able to assess the probability of meeting specific cost targets, which requires a detailed understanding of the technology; however, technical experts may be less familiar with the factors that influence costs. The task of predicting costs is as complex as forecasting technological breakthroughs, because a technology's cost depends on many interrelated factors like prices of commodities, specific manufacturing processes that are used to produce the technology, the technology's design, learning effects, and economies of scale.

Since scientists and engineers may not be the most appropriate candidates to assess these economic values, it is important to elicit additional values from economic or industry specialists who have a familiarity with specific technologies. This aligns with the general best practice of encouraging elicitations with experts from a wide range of backgrounds and viewpoints to avoid bias (Meyer and Booker, 2001; Keeney and von Winterfeldt, 1991). Another method of overcoming this limitation is to elicit only cost values from cost experts and technology performance values from technology experts. Although this would reduce the efficiency of the elicitation process, it would likely provide better quality results. Currently, there has been a tendency to elicit many values at once instead of concentrating on a few parameters, which may be negatively impacting elicitations.

The elicitation literature also suggests that it is important to have a cross-section of experts from industry, government laboratories, and academia.[4] This insight has largely been incorporated in all elicitations, though little work has been done to determine which types of experts provide the most reliable elicitation values. Preliminary research (Anadon et al., 2012) suggests that experts from industry are more pessimistic about future costs than experts in public institutions (with academics being the most optimistic).[5] There is also recognition that expert opinions may differ by country and that it is important to conduct elicitations with global experts.[6]

[4]As Morgan et al. (2009) note, selecting experts differs from the process of estimating an underlying true value through random sampling. For expert elicitations, "it is entirely possible that one expert, perhaps even one whose views are an outlier, may be correctly reflecting the underlying physical reality, and all the others may be wrong."

[5]This effect may potentially be due to a range of factors, including industry experts being most familiar with market barriers or academics' first-hand knowledge of cutting-edge technologies that are only on the brink of commercialization.

[6]The first paper to explore this issue (Anadon et al., 2012) indicates that there are significant differences between expert opinions in the US and EU.

B.3 Natural Gas Turbine Elicitations

B.3.1 Motivations

Recent advances in technologies like horizontal drilling and hydraulic fracturing have caused rapid increases in production from unconventional natural gas resources like shale formations. However, the same technologies that have facilitated this growth have also raised important questions about their environmental impacts. Natural gas is broadly considered a more environmentally benign alternative to coal due to its lower carbon dioxide emissions from combustion and its avoidance of pollutants like sulfur, particulate matter, and mercury. These environmental benefits, combined with abundant reserves, suggest that unconventional gas can play an important role in national and international energy policy—bridging a transition to a lower-carbon economy, reshaping energy security, and altering investment decisions in the electric power sector (Pacala and Socolow, 2004; DOE/EIA, 2011a).

Although abundant gas resources suggest expanded use in the electricity sector, uncertainty about the environmental impacts of production and long-run production costs makes the extent of this growth unclear (Bistline, 2012; DOE/EIA, 2011a; Coleman et al., 2011; Moniz, Jacoby, and Meggs, 2010). Additionally, natural gas price uncertainty will be influenced by the unknown policy environment, public acceptance of hydraulic fracturing (Kriesky et al., 2013; Brown et al., 2013; Brasier et al., 2011), and uncertainty surrounding life-cycle emissions (Howarth, Santoro, and Ingraffea, 2011; Jiang et al., 2011).

Another relevant uncertainty that will shape the role of natural gas in the electric power sector is the future performance of gas-turbine-based technologies. In particular, first-law efficiencies of these technologies (both with and without carbon capture) may determine the diffusion of new capacity and market share of generation from natural gas. Such characteristics are especially important for a technology subject to large fuel price volatility and to similar levelized electricity costs as other technological substitutes, which mean that even efficiency changes of a percent or two may have modest impacts on future diffusion and utilization of these technologies.

The goal of this elicitation is investigate the best practices described above through a case study of a policy-relevant technology that has been hitherto neglected in the energy technology elicitation literature. In particular, the aim of this work is to represent the current state of knowledge regarding the future of gas turbine systems for new central station electricity generation. As Table B.1 suggests, most elicitations for fossil-based electricity generation technologies have focused on coal with CCS, and when research groups look at gas with CCS, it is typically to encode uncertainty about capital costs. Here, expert judgments about the first-law efficiencies of commercially viable natural-gas-fired power plants are elicited.[7]

In the absence of this approach, most energy-economic models simply assume that future plant efficiencies will remain constant at current levels (with combined cycle efficiencies between 50 and 60 percent) or will marginally increase between now and 2050, as shown in Figure B.1. Even slight deviations from these efficiency values can have significant impacts in the development and deployment of gas-turbine-based systems, particularly when natural gas prices and climate policy are uncertain and there are many substitute technologies and fuels.

[7]Results from these expert elicitations are used as inputs to the stochastic modeling framework in Chapter 6, which assists decision-makers in the US electric power sector with capacity planning and energy technology R&D portfolio optimization under a range of technological, economic, and policy-related uncertainties.

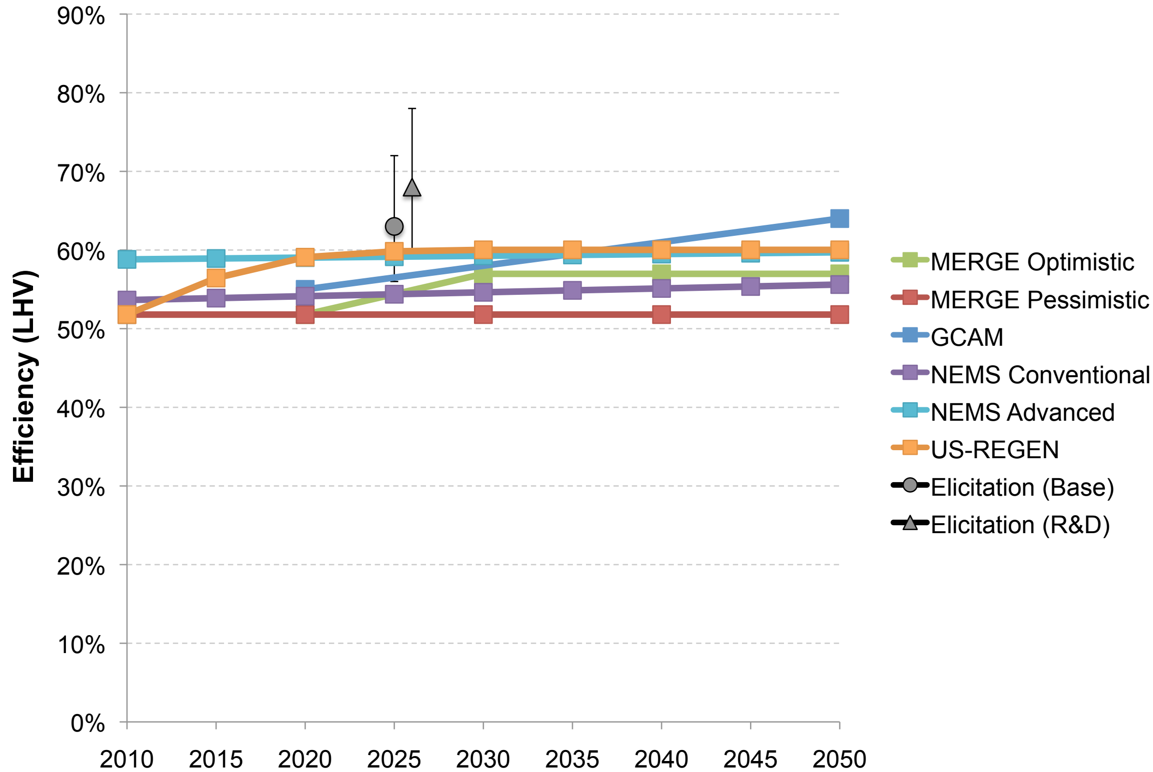

Figure B.1: First-law efficiency values (2010–2050) on a lower heating value (LHV) basis for a range of energy-economic models along with assessed range of 2025 efficiencies from this elicitation.

B.3.2 Protocol Summary

The elicitation protocol for this study was designed by drawing on the literature on techniques to minimize bias in probabilistic assessments (O'Hagan et al., 2006; Meyer and Booker, 2001; Keeney and von Winterfeldt, 1991; Morgan and Henrion, 1990; Staël von Holstein and Matheson, 1979) while addressing the specific issues raised in Section B.2.2. The protocol emphasizes robust suggestions for best practices like conducting in-person elicitations, carefully defining all terms and metrics, informing experts about common biases and strategies to avoid them (along with warm-up exercises and reminders during the elicitation discussion), and using visualization tools to facilitate quantification.

The elicitation focused on commercially viable natural-gas-fired power plants with the highest available first-law efficiency in 2025. This gas-turbine-based system should be scalable to a plant size of 500 MW and must be compliant with Clean Air Act regulations. Although the stochastic model in which this information is used contains fossil units with and without carbon capture, the elicitations considered only systems without carbon capture. These efficiency values are for commercially viable gas turbine technologies only, which is defined as having a total overnight capital cost of the system being less than or equal to $1,000 per kilowatt.[8]

The description of this plant was intentionally general to allow for the possibility that future gas-fired systems may be very different from the most commonly implemented baseload plants today, which are typically combined cycle Brayton-architecture gas turbines with bottoming steam engines. For instance, next-generation combined cycle architectures may use a gas turbine as a bottoming engine in a solid oxide fuel cell, gas turbine combination. The decision to elicit values for a single technical parameter allowed the technological experts to focus on areas within their primary domain of expertise. Restricting attention to a single value also allowed for a more in-depth discussion of how the expert viewed the history and future status of the field, which can take many hours.[9]

The second portion of the elicitation aimed to understand how enhanced public and/or private R&D programs in the US may impact the efficiencies of these technologies. There are many ways to conceptualize the success of R&D projects, as discussed in Chapter 6.1. Success can be viewed as the increased (binary) likelihood of success in reaching fixed technical or cost metrics (Baker, Chon, and Keisler, 2009a) or as an acceleration in the number of years required to reach such metrics (Blanford, 2009). The research framework here conceptualizes R&D success as adjusting the range of

[8]Expressed in terms of 2010 US dollars. This value reflects the approximate future cost of a natural gas combined cycle unit according to the Energy Information Administration's 2012 Annual Energy Outlook. The phrase "commercially viable" is used to indicate that the technology is cost-competitive with other forms of baseload electricity generation.

[9]The average elicitation session took three hours with the shortest lasting about two hours.

expected cost and performance metrics. The versatility of this probabilistic framework allows for a diverse range of representations within a stochastic programming setting, including shifting the mean of a distribution over a target R&D parameter (e.g., capital costs), reducing the variance, or eliminating fat tails (e.g., eliminating the possibility that a technology is always too expensive for deployment).

Since the selection of experts is nearly as important as the protocol itself, experts were recruited from a range of backgrounds in industry, national laboratories, and academia. Following a literature review, experts were contacted who had technological familiarity with gas-turbine-based architectures for stationary power generation with a preference for experts who could meet for in-person elicitations, who had strong technical expertise (since the focus was a technical parameter), and who are in the US. Quality control to ensure expertise was managed on the front end as assessors were being selected so that combining distributions later would not entail subjective weights. Table B.2 lists participants in the elicitations in alphabetical order.

Table B.2: List of experts and affiliations from the gas turbine elicitations.

Name	**Affiliation**
Leonard Angello	Electric Power Research Institute
Chris Edwards	Stanford University
Dale Grace	Electric Power Research Institute
Sankaran Ramakrishnan	Stanford University

Each expert received a packet in advance of the interview, which clearly defined the quantity of interest, discussed common biases, and provided a general overview of the elicitation process. The design of the elicitation protocol was based on the Stanford/SRI Assessment Protocol[10] with modifications from the literature:

1. *Motivating and Briefing*: Each session began by discussing the structure of the elicitation, by providing background about the research and how the results will be used, and by answering the expert's questions about the elicitation process.

[10]This section summarizes the primary steps and draws attention to modifications of the standard Stanford/SRI Assessment Protocol. Other authors (Staël von Holstein and Matheson, 1979; Morgan and Henrion, 1990) provide extensive information about the standard SRI Protocol.

The briefing helped experts understand the elicitation approach, to establish a sense of rapport, and to demonstrate that the elicitation was useful and worthy of serious effort.

2. *Structuring*: The next stage began by arriving at an unambiguous definition of the quantity of interest (expressed in manner that was conducive to the expert providing accurate judgments) and by determining if there were any conditioning factors that may influence the value of the quantity. This stage led into an extended technical discussion to understand how the expert saw the past, present, and future of the field. Also, this discussion allowed the experts to convey which evidence seemed most compelling and which factors and functional relationships were important for understanding the future of gas-turbine-based systems for power generation. This stage of the pre-encoding process was often the longest in the elicitation process (Staël von Holstein and Matheson, 1979).

3. *Conditioning*: The objective of this step was to condition the expert to think deeply about his or her judgment and to avoid the cognitive biases discussed in the information packet. This stage incorporated a series of warm-up questions to familiarize the expert with the concepts, structure, and techniques of the elicitation process and to get them thinking in terms of probabilities. This portion of the elicitation began with "almanac questions" for unrelated quantities and then moved to more domain-specific questions related to gas turbines.

4. *Encoding*: This stage involved the actual probability encoding process for the quantities of interest. The step began by establishing maximum and minimum credible values and by probing the expert to think carefully about these extreme values (e.g., asking for backcasts through bounding cases, where experts had to invent plausible explanations for why the true value could be lower or higher than their initial range). Once this range was chosen, cumulative probability values were elicited largely using fixed-value methods with consistency checks using fixed-probability questions. During this process, carefully articulated justifications and reasons for and against their judgments were requested.

5. *Verifying*: The objective of this final step was to test the quantitative judgments that the expert provided to ensure that the values accurately reflected their beliefs. The values given by the expert were recorded in a spreadsheet so that the results could be instantaneously plotted as both probability density functions (PDFs) and cumulative distribution functions (CDFs). Any remaining inconsistencies were resolved through conversation and iteration.

The elicited values from individual experts were later combined to summarize the current state of expert opinion in an aggregated manner. Although there are many diverse mathematical combinations and justifications for these methods (Clemen and Winkler, 1999), the linear opinion pool method was used with equal weights attached to each expert's input. There are many convenient axiomatic justifications for this approach (Clemen and Winkler, 1999) and evidence that simple combination procedures produce combined probability distributions that perform as well as those from more complicated Bayesian aggregation methods (Seaver, 1978). As mentioned before, instead of using complex calibration procedures or differential weighting, the experts in Table B.2 were selected with great care before requesting their participation and then treated all experts equally (i.e., weighting was performed up front when choosing experts instead of post-processing individual elicitation results).

Combined percentile values were fitted to shifted log-logistic distributions. These three-parameter distributions are versatile enough to represent a range of different shapes of distributions while offering a convenient way of using the 10th, 50th, and 90th percentiles to parametrize the distributions and a quantile function that is easy to use for Monte Carlo experiments. For this work, the shifted log-logistic distributions were used only as tools to visualize the PDFs and CDFs for the elicited values.

All elicitations were conducted between September and October 2012.

B.4 Results

B.4.1 Efficiency Elicitations

Figure B.2 shows the CDF of elicited values for first-law efficiencies in 2025 under the business-as-usual R&D scenario.[11] Individual values for all four experts are given along with the combined and fitted CDF. Although the figure shows some disagreement among the experts particularly for higher efficiencies, it is notable that all experts agree that the median efficiency value for 2025 will be at least 60 percent. Recall that Figure B.1 showed that only one existing energy-economic model has an efficiency value that exceeds 60 percent through 2050.[12] Thus, existing models significantly underestimate performance characteristics for future natural gas systems for electricity generation.

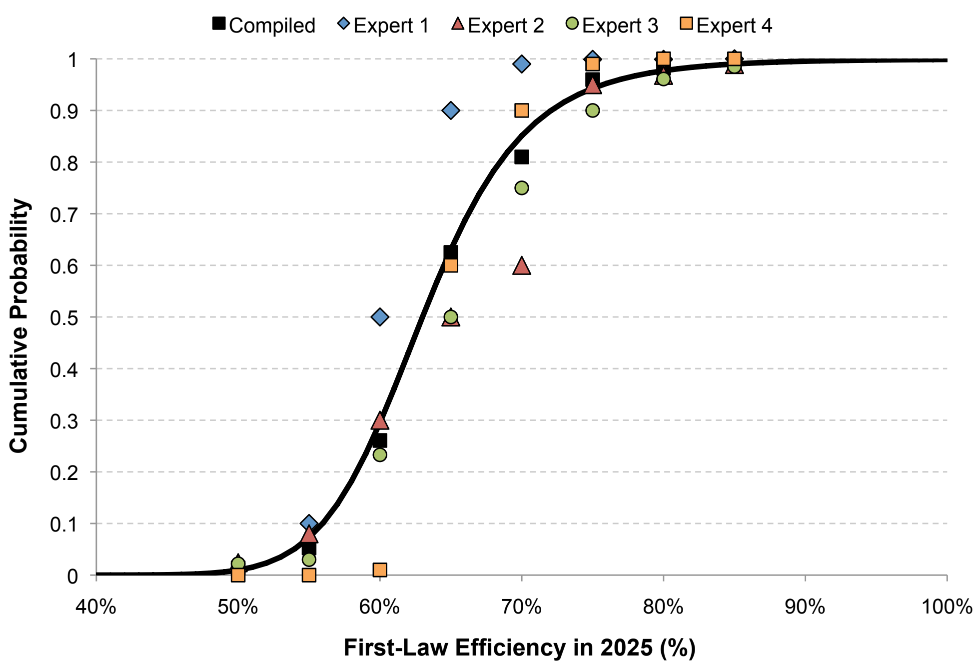

Figure B.2: Elicited values for first-law efficiencies (lower heating value basis) of gas-turbine-based electricity generators in 2025.

[11] All efficiencies for the remainder of the chapter are expressed on a lower heating value basis.

[12] The Siemens SGT5-8000H gas turbine achieved a world-record 60.75 percent efficiency in a combined-cycle configuration at the Irsching Power Station in Bavaria, Germany in May 2011.

The median first-law efficiency of the combined distribution is about 63 percent, as shown in Figure B.3. This figure compares the compiled CDFs for the business-as-usual R&D case and enhanced R&D case. These fitted values are shown as PDFs in Figure B.4. Experts believe that targeted R&D programs can increase the median efficiency from 63 to 68 percent and can increase the variance of the distribution. The increased variance suggests that the impact of research and production experience could be that new knowledge begets more uncertainty and/or opens up new possibilities for more dramatic efficiency improvements, as discussed in the next section.

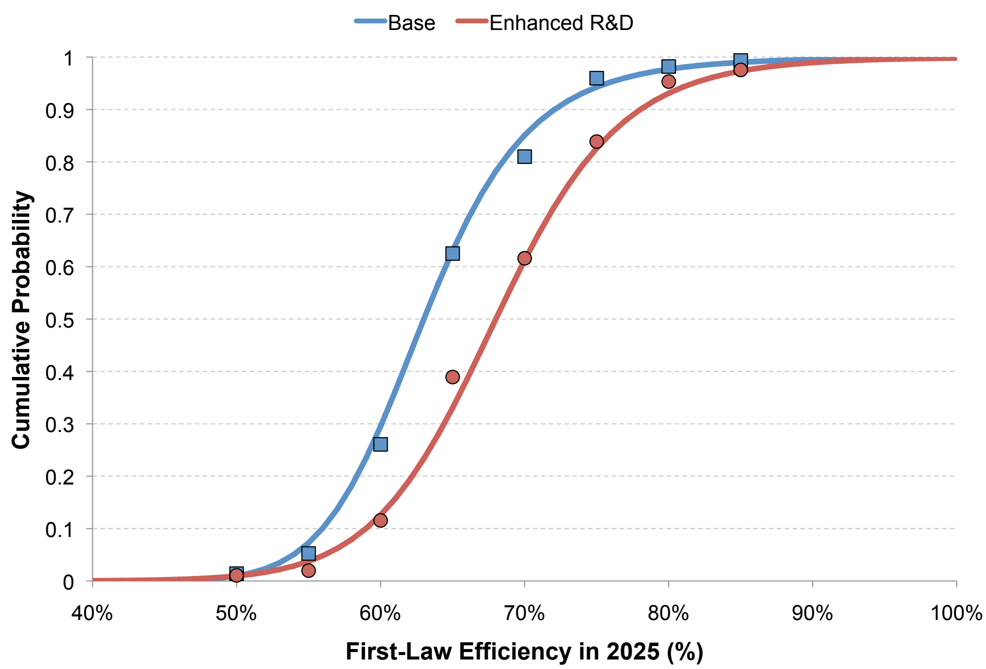

Figure B.3: Cumulative distribution functions of compiled elicitation values for the base and enhanced R&D cases.

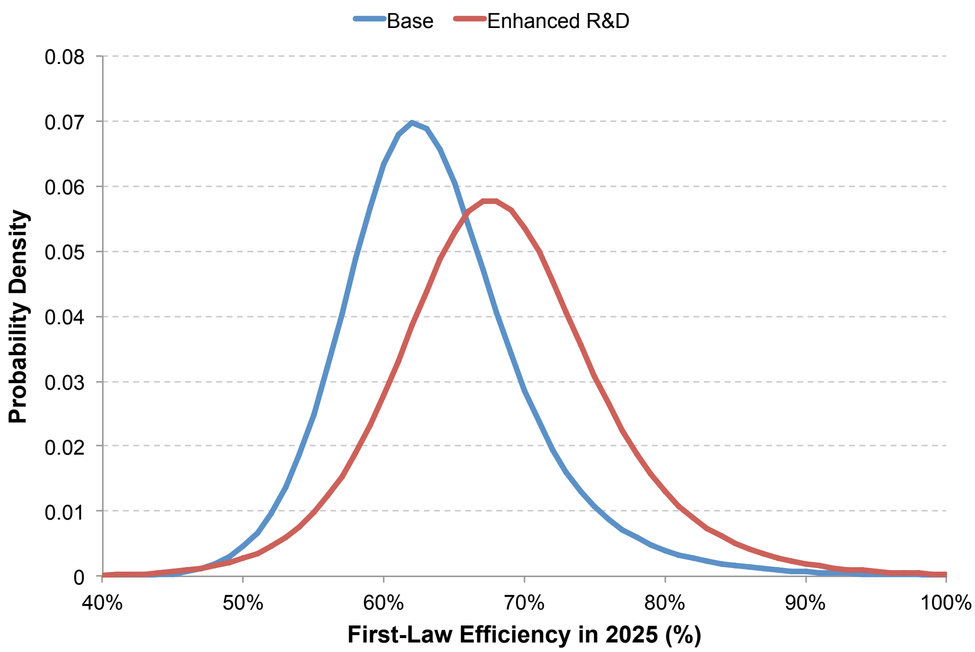

Figure B.4: Probability density functions of compiled elicitation values for the base and enhanced R&D cases.

B.4.2 Discussion

The experts agree that efficiency improvements in the coming decades will likely result from implementing existing research ideas by taking them from the laboratory, lowering costs, and implementing them at larger scales. Technological advances in gas turbine design have historically come from three sources: materials science and engineering advances, cooling improvements, and new architectures (Unger and Herzog, 1998). The lengthy technical discussions during the elicitations suggest that these factors will continue to play some part in future efficiency increases, though likely for different reasons than historical gains. When asked about prominent uncertainties that could influence the development of higher-efficiency turbine-based generators, the consensus view among the elicitations is that natural gas prices and environmental policies will play significant roles. Higher (lower) gas prices are thought to increase (lower) firms' motivation to make efficiency improvements. Experts view environmental policies (e.g., a potential federal climate policy) and regulations for emissions from existing assets (e.g., particulate matter and mercury) as important drivers for technical progress.

Progress in materials science has allowed turbine blade materials to move from conventional cast alloys in the 1960s to more highly-specialized, single-crystal alloys today (Unger and Herzog, 1998). These metallurgical advances have made high temperatures possible in combustors and turbine components. Many experts view the prospect of increasing turbine inlet temperatures and operating at higher pressure ratios as promising methods of raising efficiency values in the near term, even though efficiencies exhibit diminishing marginal returns for higher temperatures. Turbine inlet temperatures are one of the largest sources of competition between big gas turbine original equipment manufacturers (OEMs). The top priority areas for future materials research are reducing the cost of single-crystal alloys that already exist in the near term and then developing and commercializing ceramic and metal matrix composites in the longer term.[13] However, although they agree about the potential importance of ceramics, the experts disagree about the prospects for the widespread use of ceramics over the next decade.

Cooling techniques for gas turbines typically involve circulating air or steam through hot turbine components. Technological progress for cooling cascaded as a series of spillovers from military turbojet engines (where such techniques were developed in the 1960s) to civilian aircraft two to three years later, followed by diffusion to stationary power generation in approximately five years (Unger and Herzog, 1998). Many experts agree that spillovers from aerospace applications are unlikely to continue at their historical rates, as the operating profiles are very different between heavy-duty stationary gas turbines and those used for aviation (e.g., different standards for monitoring and reliability, material needs, environmental conditions, and weight restrictions). Additionally, cooling techniques advanced along with improvements in computer codes and models for finite element analysis, heat transfer, and

[13] Ceramic materials can withstand heat and corrosion and allow for higher inlet temperatures without cooling. In experimental applications as first-stage blades and combustor liners, ceramics have managed to achieve 37-degree Celsius temperature increases with associated efficiency gains of six percent (Unger and Herzog, 1998).

fluid dynamics, which were useful in modeling intricate cooling pathways, tunnels, and holes to facilitate heat transfer to the cooling fluid. The experts acknowledge that blade cooling will be an important source of temperature increases, particularly if materials science progress slows, but did not mention improved computational tools as a means of achieving these improvements.

Individual experts also suggest that first-law efficiency improvements could arise from improving auxiliary loads of the cycles themselves, from implementing more advanced architectures (e.g., intercooling, reheating, wet cycles), and from developing better heat exchangers.

The greatest disagreement between experts came in elicitations and discussions surrounding longer-term trends for gas-based architectures, especially for systems that incorporated fuel cells. Experts agree that the high end of the achievable and economic efficiency range is between 65–70 percent in the absence of dramatically new architectures. Efficiencies in this range are viewed as technically feasible but economically unlikely without enhanced R&D, which would be unlikely to come from major OEMs due to a lack of incentives for innovation or competition (outside of merely increasing inlet temperatures). The prospect of an integrated solid oxide fuel cell and gas turbine system is a highly uncertain one, though a couple of experts suggest that industry research might move toward this architecture in 10–20 years. On one hand, these systems may offer a promising route to decarbonization, since fuel cells provide an inherently high-efficiency approach to chemical separations with very high separation rates. On the other hand, such systems are currently only demonstrable at a laboratory level and would face numerous hurdles to commercialization due to concerns about the overall economics of the system, the longevity of the fuel cell, the stability of the membranes, and the ability to increase the packing density and decreasing size by a factor of ten. Experts disagree about the likelihood of achieving the required performance and cost targets for this fuel cell system even with targeted R&D. This sense of uncertainty about advanced turbine-based architectures and technical progress in the mid- to long-term future accounts for the large variance for the enhanced R&D distribution in Figure B.3.

As mentioned at the beginning of this section, it is not clear *prima facie* whether future performance and cost trends for turbine-based electricity generators will follow historical values. Although there are many promising developments on the horizon, there are also many reasons to doubt that historical sources of technological change (e.g., spillovers from the aerospace industry, rapid advances in computational fluid dynamics, or increasing turbine inlet temperatures) will continue to be primary drivers of efficiency gains in the future. Consequently, expert elicitations fill this void by providing a basis for forecasting future efficiency values for gas-based systems.

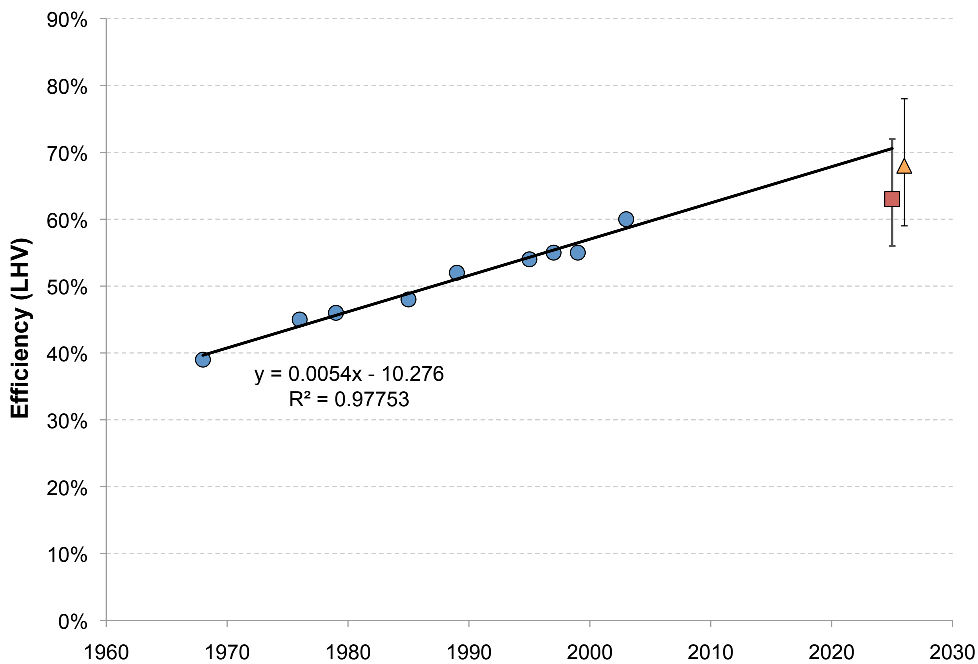

Figure B.5: Historical values of best-available combined-cycle efficiencies (1968–2003) with a linear trendline. The values at 2025 represent the median combined expert elicitation values for the base R&D (red square) and enhanced R&D (orange triangle) with the 10th and 90th percentiles shown with error bars.

Figure B.5 shows the historical values for combined-cycle efficiencies in the US electric power sector between 1968–2003. A simple linear trendline, when extrapolated to 2025, suggests that efficiencies would reach upward of 70 percent. Although this efficiency falls within the 10th and 90th percentiles of the elicitation values, the median estimates under business-as-usual and enhanced R&D conditions are notably lower than this trendline. Thus, the expert elicitations support the conclusion that prospective efficiency gains are unlikely to follow historical trends. However, these median values are still appreciably higher than the efficiencies used in many integrated assessment models.

B.5 Summary and Extensions

In addition to the insights about the future of gas turbine systems discussed in the previous section, these elicitations illustrated many best practices for conducting expert elicitations.

The largest takeaway was that face-to-face elicitations are extremely useful in critically examining experts' reported probability values, particularly for the tails of the distribution. Feedback questions for participants' responses make them think critically about the values they give and force them to brainstorm how extreme values may be lower or higher than their initial impressions suggest. In one elicitation, a question was reframed in three different ways before the expert noted the possibility of using supercritical water injection in the combustor and revised the efficiency estimate upward. During the debriefing sessions, subjects reported discomfort in thinking about tail probabilities and suggested that, without the interviewer's intervention, they would have selected an anchor value and then extrapolated to select other values. Additionally, the warm-up exercises suggested that the experts were initially overconfident, as the actual number of "surprises" (i.e., values falling outside of the 10^{th} and 90^{th} percentiles) was over twice as high as the expected number of surprises in three of four cases. Thus, based on these observations, future research should examine to what degree at-a-distance elicitations exhibit greater overconfidence compared with in-person protocols and how interactive digital tools can bridge this gap if it exists.

Many other advantages of conducting in-person elicitations were observed:

- In-person elicitations allow the interviewer to clarify misconceptions that may not be noticed without asking probing questions. This technique was invoked to determine whether an expert was conditioning on events that were not discussed, to clarify specific instances of how experts can avoid biases during the actual elicitation, and to resolve a misunderstanding about the definition of cumulative probabilities, which was discovered when the interviewer noticed and inconsistency in the given values.

- Conducting an in-person elicitation indicates that the interviewer cares about the quality of the elicitations and the results of the assessment.

- Many subjects reported that they were more comfortable eliciting the values face-to-face due to the ability to ask the interviewer questions.

Ultimately, one of the largest benefits of the elicitation process is that it gives modelers more opportunities to consult technical experts who have the greatest experience and familiarity with technologies. These experts also have knowledge that energy-economic models may not capture but is important to the development and deployment of technologies. Since these insights typically come out in unstructured conversation, at-a-distance elicitations bypass (or do not take full advantage of) these deep interactions. This point also implies that elicitations have an important role in energy modeling even in a deterministic setting. For instance, exogenous technological progress in deterministic models is typically informed by engineering cost estimates, which should rely on elicitations to assess expert opinion and to structure sensitivities. No matter the model structure, elicitations can help modelers to identify and avoid potential blind spots in the planning process. This function is particularly salient for energy modeling in the context of climate change, which prominently features a few nascent technologies.[14]

Expert elicitations are as important for future modeling efforts as they are for those in the present. Probabilistic assessments preserve information about current beliefs for use in the future, which means that formally capturing such beliefs is necessary

[14]For instance, elicitations in Baker, Chon, and Keisler (2009b) suggest that the prospects of technical success for post-combustion carbon capture technologies are still controversial among experts in the area, even though many energy models take the availability of such technologies as given.

for hindcasting exercises. Therefore, elicitations play an integral part in constructing information management systems, improving models for decision support, and combating hindsight bias. These assessments are likewise necessary for evaluating the dynamics of learning (Hannart, Ghil, and Dufresne, 2013; Oppenheimer, O'Neill, and Webster, 2008) and for understanding why errant forecasts were wrong (Craig, Gadgil, and Koomey, 2002). Modelers should compare forecasts with evolving observations to determine trends in estimation errors and to diagnose any systematic forecast biases.

A stochastic analysis is only as good as the probability encoding process behind it. The usefulness of models and elicitation processes would be enhanced if future research compared face-to-face, online, phone, and written elicitations. There are currently no empirical assessments of whether there is an upward or downward bias to moments of distributions based on whether elicitations are conducted in person or at a distance, though the experience here suggests that at-a-distance methods likely underestimate tail probabilities. These experiments could explore how interactive digital elicitation tools can bridge the gap between in-person elicitations (which are recommended by decision analysis practitioners) and at-a-distance paper elicitations (which are prevalent due to their cost-effectiveness and economies of scale). Answers to these questions are especially relevant given the need for more frequent elicitations involving rapidly changing technologies like solar (Reichelstein and Yorston, 2013), where it is important to use techniques that can save time and money while not compromising quality.

Appendix C

Comparing R&D Success Valuations

C.1 Motivations and Problem Description

The research and development (R&D) literature, particularly for energy technologies, has primarily concentrated on the probabilistic relationship between R&D investments and potential outcomes while treating the valuation of these outcomes deterministically. It is common to represent R&D benefits through standard energy-economic or integrated assessment models, which typically do not represent uncertainty explicitly and do not capture the hedging potential and optionality of technologies. These models value R&D by using the wait-and-see (learn-then-act) approach, which assumes that uncertainty about exogenous market conditions is resolved at the beginning of the time horizon. This solution suggests that perfect information is available before capacity planning decisions are made.

Such information is usually not available when early-stage deployment and R&D decisions are made, and the decision-maker must select a strategy that hedges against possible contingencies in an uncertain environment. This here-and-now approach seeks to identify a stochastic solution, which is the same under all states of the world until uncertainties are resolved and recourse decisions can be made. The explicit inclusion of uncertainty is critical in energy-related decisions, because exogenous market uncertainties are pervasive and may be key drivers of diffusion. Such uncertainties may include advances in substitute/complementary/enabling technologies, demand, public acceptance, and the regulatory environment.

The stylized example presented here demonstrates how the value of R&D success can vary based on the decision-making approach used for capacity planning and dispatch decisions. This example illustrates how the influence of different approaches on R&D valuation is ambiguous so that there is no necessary relationship between the value of R&D success and the decision-making approach used to evaluate program benefits.[1] The success valuation depends on characteristics of the decision problem (e.g., the objective function formulation, constraints, and parameterization), the uncertainties considered, and the nature of the R&D program.

Comparing the stochastic hedging (i.e., here-and-now) valuation with the perfect information (i.e., wait-and-see) approach, the necessary and sufficient condition for a higher value of R&D success under the wait-and-see approach is:

$$z^*(\theta) - z^*(\theta') \leq z_{ws}(\theta) - z_{ws}(\theta')$$
$$\min_x \mathbb{E}_\omega f(x;\omega,\theta) - \min_x \mathbb{E}_\omega f(x;\omega,\theta') \leq \mathbb{E}_\omega \left[\min_x f(x;\omega,\theta) - \min_x f(x;\omega,\theta')\right]$$
$$\mathrm{EVPI}(\theta) \leq \mathrm{EVPI}(\theta')$$

where θ is the baseline technological state (distribution) without R&D, and θ' is the advanced technological state (distribution) with R&D success. Since the direction of the inequality depends on characteristics of the uncertainties and optimization problem, the relationship between the decision-making approach and the R&D success valuation is ambiguous.

For analytical clarity, this example compares the wait-and-see and expected-value approaches for R&D success valuation. The expected value of the here-and-now (stochastic) solution lies between these two solutions.

This example is a simplified representation of the capacity planning problem, where a utility decides which type of power plant to build to meet demand. The choice set consists of two alternatives: a conventional carbon-intensive generator (e.g.,

[1]Mathematical definitions of these solution approaches are found in Chapter 3.3.2.

coal) and a low-carbon unit (e.g., solar).

The decision is complicated by an uncertain carbon tax that does not materialize until the second stage. If a fossil-fueled generator is built in the first stage, the utility bears the risk that the carbon tax will make operations in the second stage prohibitively expensive, which could require the construction of a low-carbon unit. If a low-carbon unit is built in the first stage, the utility risks the possibility that no climate policy will be implemented and that the larger upfront investment cost will have been unnecessary.

The utility also must decide on the timing of this investment given the possibilities of delaying the capacity installation decision and of R&D reducing the second-stage investment costs of the two alternatives. If R&D is expected to lower capital costs in the second stage, it might be preferable to delay the decision to take advantage of these lower costs. This decision depends on the assumed cost of delay and on expectations about the efficacy of R&D in changing cost outcomes. Without loss of generality, this example assumes the R&D outcome is deterministic so that the decision-maker knows at the beginning of the first stage whether the program will be successful and how R&D will influence investment costs.

C.2 Notation

The decision variables and parameters for the utility's optimization problem use the following sets and corresponding index notation:

Sets and Indices

$i \in \mathcal{I}$ generation technology type; $i \in \{1, 2\}$

$t \in \mathcal{T}$ stage; $t \in \{1, 2\}$

$\omega \in \Omega$ carbon policy scenario; $\omega \in \{1, 2\}$

Generator 1 is the conventional fossil-fueled unit, and generator 2 is the advanced low-carbon unit. The first stage corresponds to near-term decisions over the next decade or so, and the second stage represents subsequent decisions. It is assumed that the carbon tax uncertainty is not resolved until the second stage. Scenario 1 ($\omega = 1$)

corresponds to the state of the world where there is no climate policy, and scenario 2 ($\omega = 2$) corresponds to the carbon tax state. If an R&D program is successful, the cost reductions do not take effect until the second stage.

Decision Variables

x_i	first-stage investment in generation technology i
y_i^ω	second-stage investment of technology i in scenario ω
w_i^ω	second-stage dispatch of technology i in scenario ω
d	unmet demand in first stage

All decision variables are binary so that 1 indicates investment (or dispatch) and 0 indicates no investment.

Parameters

c_i	first-stage investment cost of technology i
f_i	second-stage investment cost of technology i with no R&D
e_i	second-stage investment cost of technology i with an R&D success
Δ	cost of lost demand
δ	discount factor at time t
$\widetilde{g}_i$	uncertain second-stage operating costs of technology i
p	probability of no carbon tax ($\omega = 1$)
$1 - p$	probability of carbon tax ($\omega = 2$)

Operating costs during the second stage ($\widetilde{g}_i^\omega$) are random variables due to an uncertain climate policy. The low-carbon technology is assumed to have negligible operating costs. Demand must be met through generation in the second stage but can be satisfied in the first stage either through generation or by paying a cost Δ.

R&D lowers the investment cost of technology i during the second stage from f_i to e_i. The R&D program is deterministic in the sense that the decision-maker knows at the beginning of the time horizon whether the second-stage investment cost is f_i or e_i.

C.3 Approaches

C.3.1 Expected-Value Approach

$$\hat{z}_d = \min\ f(x, \bar{\omega}) = \min\ c_1 x_1 + c_2 x_2 + \Delta d + \delta \left[f_1 y_1 + f_2 y_2 + \left(\sum_{\omega \in \Omega} p(\omega) g_1^{\omega} \right) w_1 \right] \tag{C.1}$$

$$\text{s.t. } x_1 + x_2 + d \geq 1 \tag{C.2}$$

$$x_1 + x_2 + y_1 + y_2 \geq 1 \tag{C.3}$$

$$x_1 + y_1 = w_1 + y_2 \tag{C.4}$$

Equation C.1 represents the objective function, which minimizes the discounted sum of investment and operating costs for the first and second stages. Equations C.2 and C.3 signify the first- and second-stage load balances, respectively. Equation C.4 represents the dispatch constraint for the fossil-fueled generation technology. The expected-value solution is $x_d \in \text{argmin}\ \{f(x, \bar{\omega}) \mid x \in C^{\bar{\omega}}\}$, and the expected value of the expected-value solution is $z_d = \mathbb{E}\left[f(x_d, \omega)\right]$.

C.3.2 Wait-and-See Approach

$$z^{\omega} = \min\ f(x, \omega) = \min\ c_1 x_1 + c_2 x_2 + \Delta d + \delta \left[f_1 y_1 + f_2 y_2 + g_1^{\omega} w_1 \right]$$

$$\text{s.t. } x_1 + x_2 + d \geq 1$$

$$x_1 + x_2 + y_1 + y_2 \geq 1$$

$$x_1 + y_1 = w_1 + y_2$$

The expected value of the wait-and-see solution is $z_{ws} = \mathbb{E}\left[\min\ f(x, \omega)\right]$.

C.4 Cases

C.4.1 Case 1: $i = 2$ and large Δ

This case involves an R&D program aimed at reducing the cost of the low-carbon technology (e.g., solar) and has a prohibitively large cost of delay. This high cost means that the decision-maker cannot postpone an installation decision until the second stage when more information is known about the climate policy.

Proposition C.4.1 (Expected-Value Approach) *A parameterization exists such that the optimal expected-value solution for both the f_2 and e_2 cases (i.e., regardless of R&D success for the low-carbon technology) is:* $x_1 = 1, x_2 = d = 0; y_1^1 = y_2^1 = y_1^2 = 0, y_2^2 = 1, w_1^1 = 1, w_1^2 = 0.$

Proof For the expected-value strategy to build a conventional generator during the first stage ($x_1 = 1$), the projected cost of building and operating the facility must be less than or equal to the investment cost of the advanced option: $c_1 + \delta\left[pg_1^1 + (1-p)g_1^2\right] \leq c_2$. The other condition required to ensure that a conventional unit is built in the first stage instead of delaying investment until the second stage is: $c_1 + \delta\left[pg_1^1 + (1-p)g_1^2\right] \leq \Delta + \delta \min\left[f_1 + pg_1^1 + (1-p)g_1^2,\ f_2\right]$. In this case, the cost of delay (Δ) is large, which suggests that this inequality will always be satisfied. ∎

Proposition C.4.2 (Wait-and-See Approach) *For both the f_2 and e_2 cases, a parameterization exists such that the optimal wait-and-see solution under the no-policy scenario ($\omega = 1$) is $x_1 = w_1 = 1, x_2 = d = y_1 = y_2 = 0$ and that the solution under the policy scenario ($\omega = 2$) is $x_2 = 1, x_1 = d = y_1 = y_2 = w_1 = 0$.*

Proof For $\omega = 1$, the necessary conditions for a conventional unit to be build are $c_1 + \delta g_1^1 \leq c_2$ and $c_1 + \delta g_1^1 \leq \Delta + \delta(e_1 + g_1^1)$. For $\omega = 2$, the necessary inequalities to ensure that the advanced, low-carbon generator is built are $c_2 \leq c_1 + \delta g_1^2$ and $c_2 \leq \delta(\Delta + e_2)$. Since Δ is assumed to be sufficiently large in this case, these necessary conditions can be summarized as: $\delta g_1^1 \leq c_2 - c_1 \leq \delta g_1^2$. ∎

Proposition C.4.3 *There exists a parameterization for the capacity planning problem such that the expected-value and wait-and-see solutions above hold simultaneously.*

Proof The necessary condition from the proof of Proposition C.4.1 can be rewritten as $\delta\left[pg_1^1 + (1-p)g_1^2\right] \leq c_2 - c_1$. The left-hand side is greater than δg_1^1, since $\delta\left[pg_1^1 + (1-p)g_1^2\right]$ is a convex combination of δg_1^1 and δg_1^2 (where $g_1^2 > g_1^1$). Thus, the necessary conditions for the two decision-making approaches are compatible if $\boxed{\delta\left[pg_1^1 + (1-p)g_1^2\right] \leq c_2 - c_1 \leq \delta g_1^2}$. $\blacksquare$

The value of R&D success (i.e., benefit from reduced costs) for the expected-value approach can be expressed as:

$$\begin{aligned} z_d(f) - z_d(e) &= \left[p(c_1 + \delta g_1^1) + (1-p)(c_1 + \delta f_2)\right] - \left[p(c_1 + \delta g_1^1) + (1-p)(c_1 + \delta e_2)\right] \\ &= \delta(1-p)(f_2 - e_2) > 0 \end{aligned}$$

The value of R&D success for the wait-and-see approach can be expressed as:

$$\begin{aligned} z_{ws}(f) - z_{ws}(e) &= \left[p(c_1 + \delta g_1^1) + (1-p)c_2\right] - \left[p(c_1 + \delta g_1^1) + (1-p)c_2\right] \\ &= 0 \end{aligned}$$

Therefore, in the case where the opportunity cost of delay is large and R&D applies to the low-carbon technology, the value of R&D is larger for the expected-value approach than for the wait-and-see approach. The objective function value under both R&D conditions is smaller for the wait-and-see approach, but this gap is smaller in the scenario when the R&D program is successful. By the time that the lower-cost technology is available in the second stage, the only approach that can take advantage of this cost reduction is the expected-value one when the carbon tax scenario is realized. The expected-value strategy builds the low-carbon generator in the second stage due to the strategy's irreversible investment in the carbon-intensive generator during the first stage, which is prohibitively costly when a carbon tax is in place. The substantial cost of delay means that the wait-and-see approach always invests in the first stage, and the expected-value strategy only invests during the second stage when the first-stage decision is suboptimal (i.e., when the carbon tax is much more stringent than the expected value).

C.4.2 Case 2: $i = 1$ and small Δ

This scenario represents the case with an R&D program for the conventional technology and small cost of delay. This small cost can be interpreted as a utility's ability to engage in power purchase agreements at relatively low costs in the near term.

Proposition C.4.4 (Expected-Value Approach) *A parameterization exists such that the optimal expected-value solution for both the f_2 and e_2 cases is: $x_2 = 1, x_1 = d = 0; y_i^\omega = w_1^\omega = 0$.*

Proof For the expected-value strategy to build a low-carbon generator during the first stage ($x_2 = 1$), the projected cost of building the unit must be less than the investment and operating cost of the conventional option: $c_2 \leq c_1 + \delta\left[pg_1^1 + (1-p)g_1^2\right]$. The other condition required to ensure that an advanced, low-carbon unit is built in the first stage instead of delaying investment until the second stage is: $c_2 \leq \Delta + \delta\left[e_1 + pg_1^1 + (1-p)g_1^2\right]$. When Δ is small, it can be assumed that $\Delta + \delta e_1 \leq c_1$, which means that the second condition is more restrictive. ∎

Proposition C.4.5 (Wait-and-See Approach) *For the no R&D case, a parameterization exists such that the optimal wait-and-see solution is $x_1 = 1$ when $\omega = 1$ and $x_2 = 1$ when $\omega = 2$. In the R&D case for the conventional technology, the solution is $d = y_1 = 1$ when $\omega = 1$ and $x_2 = 1$ when $\omega = 2$.*

Proof For the no R&D case, the necessary condition for the optimality of the above solution is: $p(c_1 + \delta g_1^1) + (1-p)c_2 \leq p\left[\Delta + \delta(f_1 + g_1^1)\right] + (1-p)c_2$. This condition simplifies to $c_1 \leq \Delta + \delta f_1$, which is always true if $\delta = 1$. For the R&D case, the necessary condition to ensure that delay is optimal under the no-policy scenario instead of building the conventional generator in the first stage is: $p\left[\Delta + \delta(e_1 + g_1^1)\right] + (1-p)c_2 \leq p(c_1 + \delta g_1^1) + (1-p)c_2$, which can be reduced to $\Delta + \delta e_1 \leq c_1$. ∎

Proposition C.4.6 *There exists a parameterization for the capacity planning problem such that the expected-value and wait-and-see solutions above hold simultaneously.*

Proof The necessary condition from Proposition C.4.4 can be expressed in the form $c_2 - \delta\left[pg_1^1 + (1-p)g_1^2\right] \leq \Delta + \delta e_1$. Thus, combining this result with the condition from Proposition C.4.5, the necessary conditions for the two decision-making approaches are compatible if $\boxed{c_2 - \delta\left[pg_1^1 + (1-p)g_1^2\right] \leq \Delta + \delta e_1 \leq c_1}$. The lower bound ensures that the expected-value solution will choose the low-carbon generator during the first stage and not opt to delay in the R&D case.[2] The upper bound ensures that, in the no-policy scenario, the wait-and-see solution will wait to build a conventional generator during the second stage when R&D has reduced its cost, which is true if Δ is small. ∎

The value of R&D success for the expected-value approach can be expressed as:

$$z_d(f) - z_d(e) = [pc_2 + (1-p)c_2] - [pc_2 + (1-p)c_2] = 0$$

The value of R&D success for the wait-and-see approach can be expressed as:

$$\begin{aligned} z_{ws}(f) - z_{ws}(e) &= \left[p(c_1 + \delta g_1^1) + (1-p)c_2\right] - \left[p(\Delta + \delta(e_1 + g_1^1)) + (1-p)c_2\right] \\ &= p(c_1 - \Delta - \delta e_1) > 0 \end{aligned}$$

Therefore, in the case where the opportunity cost of delay is small and R&D applies to the conventional technology, the value of R&D is larger for the wait-and-see approach than for the expected-value approach.

C.5 Summary of Findings

This example shows how it is ambiguous as to whether using different decision-making approaches will increase or decrease the value of R&D successes. Consequently, the direction of optimal investments in R&D programs is ambiguous as well. The impact of the decision-making approach on R&D valuation depends on factors like:

[2]This condition is true if $c_2 \approx e_1$ (i.e., the investment cost of the low-carbon unit is relatively small), p is small (i.e., the probability of a carbon tax is high), and/or g_1^2 is large (i.e., the carbon tax is stringent).

- How uncertainties interact in the decision problem (i.e., objective function, constraints, and parameterization of the optimization problem)
- The form of the distributions chosen for exogenous uncertainties
- The change in technological characteristics brought about by successful R&D

Thus, it requires modeling efforts like those in Chapter 6 to determine the influence of decision-making approaches on R&D success valuations in specific contexts.

Bibliography

Abdulla, A., I. L. Azevedo, and M. G. Morgan. 2013. "Expert Assessments of the Cost of Light Water Small Modular Reactors." *Proceedings of the National Academy of Sciences of the United States of America* (Forthcoming).

Ackerman, F., E. A. Stanton, and R. Bueno. 2010. "Fat Tails, Exponents, Extreme Uncertainty: Simulating Catastrophe in DICE." *Ecological Economics* 69 (8):1657–1665.

AEC. 1969. "Reactor Operating Experiences." Tech. Rep. 69-10, US Atomic Energy Commission.

Alic, J. A., D. C. Mowery, and E. S. Rubin. 2003. "US Technology and Innovation Policies: Lessons for Climate Change." Pew Center on Global Climate Change.

Allen, D. T., V. M. Torres, J. Thomas, D. W. Sullivan, M. Harrison, A. Hendler, S. C. Herndon, C. E. Kolb, M. P. Fraser, A. D. Hill, B. K. Lamb, J. Miskimins, R. F. Sawyer, and J. H. Seinfeld. 2013. "Measurements of Methane Emissions at Natural Gas Production Sites in the United States." *Proceedings of the National Academy of Sciences of the United States of America* 110 (44):17768–17773.

Allen, M. R. and D. J. Frame. 2007. "Call Off the Quest." *Science* 318 (5850):582–583.

Anadon, L. D., V. Bosetti, M. Bunn, M. Catenacci, and A. Lee. 2012. "Expert Judgments about RD&D and the Future of Nuclear Energy." *Environmental Science and Technology* 46 (21):11497–11504.

Anadon, L. D., M. Bunn, G. Chan, M. Chan, C. Jones, R. Kempener, A. Lee, N. Logar, and V. Narayanamurti. 2011. "Transforming US Energy Innovation." Harvard Kennedy School, Belfer Center for Science and International Affairs.

Anthoff, D., R. S. J. Tol, and G. W. Yohe. 2009. "Risk Aversion, Time Preference, and the Social Cost of Carbon." *Environmental Research Letters* 4 (2).

Arrow, K. J. and A. C. Fischer. 1974. "Environmental Preservation, Uncertainty, and Irreversibility." *The Quarterly Journal of Economics* 88 (2):312–319.

Atkinson, G. D., S. Dietz, J. Helgeson, C. J. Hepburn, and H. Saelen. 2009. "Siblings, Not Triplets: Social Preferences for Risk, Inequality, and Time in Discounting Climate Change." *Economics* 3.

Ayyub, B. M. 2001. *Elicitation of Expert Opinions for Uncertainty and Risks*. Boca Raton, FL: CRC Press.

Baker, E. 2009. "Optimal Policy under Uncertainty and Learning about Climate Change: A Stochastic Dominance Approach." *Journal of Public Economic Theory* 11 (5):721–747.

Baker, E., H. Chon, and J. Keisler. 2008. "Advanced Nuclear Power: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy." University of Massachusetts Amherst.

———. 2009a. "Advanced Solar R&D: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy." *Energy Economics* 31 (Supplement 1):S37–S49.

———. 2009b. "Carbon Capture and Storage: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy." *Climatic Change* 96 (3):379–408.

———. 2010. "Battery Technology for Electric and Hybrid Vehicles: Combining Expert Views about Prospects for Advancement." *Technological Forecasting and Social Change* 77 (7):1139–1146.

Baker, E. and J. Keisler. 2011. "Cellulosic Biofuels: Expert Views on Prospects for Advancement." *Energy* 36 (1):595–605.

Baker, E. and E. Shittu. 2008. "Uncertainty and Endogenous Technical Change in Climate Policy Models." *Energy Economics* 30 (6):2817–2828.

Baker, E. and S. Solak. 2011. "Climate Change and Optimal Energy Technology R&D Policy." *European Journal of Operational Research* 213 (2):442–454.

———. 2013. "Management of Energy Technology for Sustainability: How to Fund Energy Technology Research and Development." *Production and Operations Management* (Forthcoming).

Baker, M. B. and G. H. Roe. 2009. "The Shape of Things to Come: Why Is Climate Change So Predictable?" *Journal of Climate* 22 (17):4574–4589.

Barbose, G., R. Wiser, A. Phadke, and C. Goldman. 2008. "Reading the Tea Leaves: How Utilities in the West Are Managing Carbon Regulatory Risk in Their Resource Plans." Tech. Rep. LBNL-44E, Lawrence Berkeley National Laboratory: Environmental Energy Technologies Division.

Beale, E. M. L. 1955. "On Minimizing a Convex Function Subject to Linear Inequalities." *Journal of the Royal Statistical Society* 17 (2):173–184.

Bier, V. M., Y. Y. Haimes, J. H. Lambert, N. C. Matalas, and R. Zimmerman. 1999. "A Survey of Approaches for Assessing and Managing the Risk of Extremes." *Risk Analysis* 19 (1):83–94.

Birge, J. R. 1982. "The Value of the Stochastic Solution in Stochastic Linear Programs with Fixed Recourse." *Mathematical Programming* 24 (1):314–325.

———. 1995. "Models and Model Value in Stochastic Programming." *Annals of Operations Research* 59 (1):1–18.

Birge, J. R. and F. Louveaux. 2011. *Introduction to Stochastic Programming*. New York, NY: Springer.

Birge, J. R. and C. H. Rosa. 1996. "Incorporating Investment Uncertainty into Greenhouse Policy Models." *The Energy Journal* 17 (1):79–90.

Bistline, J. E. 2012. "Electric Sector Capacity Planning under Uncertainty: Shale Gas and Climate Policy in the US." In *31st USAEE/IAEE North American Conference*. Austin, TX.

Bistline, J. E. and J. P. Weyant. 2013. "Electric Sector Investments under Technological and Policy-Related Uncertainties: A Stochastic Programming Approach." *Climatic Change* 121 (2):143–160.

Blanford, G. J. 2005. "A Global Portfolio Strategy for Climate Change Technology Development." In *International Energy Workshop*. Kyoto, Japan.

———. 2006. *Technology Strategy for Climate Change: Optimal R&D Investment and Its Interaction with Abatement Policy*. Ph.D. thesis, Stanford University.

———. 2009. "R&D Investment Strategy for Climate Change." *Energy Economics* 31 (Supplement 1):S27–S36.

———. 2013. "Is Overshoot Optimal? Negative Emissions in a Cost-Benefit Setting." In *International Energy Workshop*. Paris, France.

Blanford, G. J. and J. P. Weyant. 2007. "Optimal Investment Portfolios for Basic R&D." Working Paper. Stanford University.

Bloom, N., M. Schankerman, and J. Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4):1347–1393.

Bosetti, V., C. Carraro, A. Sgobbi, and M. Tavoni. 2009. "Delayed Action and Uncertain Stabilisation Targets. How Much Will the Delay Cost?" *Climatic Change* 96 (3):299–312.

Bosetti, V., M. Catenacci, G. Fiorese, and E. Verdolini. 2012. "The Future Prospect of PV and CSP Solar Technologies: An Expert Elicitation Survey." *Energy Policy* 49:308–317.

Bosetti, V. and M. Tavoni. 2009. "Uncertain R&D, Backstop Technology, and GHGs Stabilization." *Energy Economics* 31 (Supplement 1):S18–S26.

Botterud, A., M. D. Ilic, and I. Wangensteen. 2005. "Optimal Investments in Power Generation under Centralized and Decentralized Decision Making." *IEEE Transactions on Power Systems* 20 (1):254–263.

Bradbury, J., M. Obeiter, L. Draucker, W. Wang, and A. Stevens. 2013. "Clearing the Air: Reducing Upstream Emissions from US Natural Gas Systems." World Resources Institute.

Brasier, K. J., M. R. Filteau, D. K. McLaughlin, J. Jacquet, R. C. Stedman, T. W. Kelsey, and S. J. Goetz. 2011. "Residents' Perceptions of Community and Environmental Impacts from Development of Natural Gas in the Marcellus Shale: A Comparison of Pennsylvania and New York Cases." *Journal of Rural Social Sciences* 26 (1):32–61.

Brown, E., K. Hartman, C. P. Borick, B. G. Rabe, and T. M. Ivacko. 2013. "The National Surveys on Energy and Environment Public Opinion on Fracking: Perspectives from Michigan and Pennsylvania." Center for Local, State, and Urban Policy (CLOSUP).

Brysse, K., N. Oreskes, J. O'Reilly, and M. Oppenheimer. 2013. "Climate Change Prediction: Erring on the Side of Least Drama?" *Global Environmental Change* 23 (1):327–337.

Burnham, A., J. Han, C. E. Clark, M. Wang, J. B. Dunn, and I. Palou-Rivera. 2012. "Life-Cycle Greenhouse Gas Emissions of Shale Gas, Natural Gas, Coal, and Petroleum." *Environmental Science and Technology* 46 (2):619–627.

Butler, M. P., P. M. Reed, K. Fisher-Vanden, K. Keller, and T. Wagener. 2013. "Identifying Parametric Controls and Dependencies in Integrated Assessment Models Using Global Sensitivity Analysis." *Environmental Modelling and Software* (Under Revision).

Cathles, L. M., L. Brown, and M. Taam. 2012. "A Commentary on 'The Greenhouse-Gas Footprint of Natural Gas in Shale Formations' by R. W. Howarth, R. Santoro, and Anthony Ingraffea." *Climatic Change* 113 (2):525–535.

Chan, G., L. D. Anadon, M. Chan, and A. Lee. 2011. "Expert Elicitation of Cost, Performance, and RD&D Budgets for Coal Power with CCS." *Energy Procedia* 4:2685–2692.

Clemen, R. T. and R. L. Winkler. 1999. "Combining Probability Distributions from Experts in Risk Analysis." *Risk Analysis* 19 (2):187–203.

Cohen, L. R. and R. G. Noll. 1991. *The Technology Pork Barrel*. Washington, DC: The Brookings Institution.

Coleman, J. L., R. C. Milici, T. A. Cook, R. R. Charpentier, M. Kirschbaum, T. R. Klett, R. M. Pollastro, and C. J. Schenk. 2011. "Assessment of Undiscovered Oil and Gas Resources of the Devonian Marcellus Shale of the Appalachian Basin Province, 2011." Tech. Rep. US Geological Survey Fact Sheet 20113092, US Geological Survey.

Conrad, J. M. 1980. "Quasi-Option Value and the Expected Value of Information." *The Quarterly Journal of Economics* 94 (4):813–820.

Cooke, R. M. 2013. "Uncertainty Analysis Comes to Integrated Assessment Models for Climate Change. . . and Conversely." *Climatic Change* 117 (3):467–479.

Cooke, R. M., D. Nieboer, and J. Misiewicz. 2011. "Fat-Tailed Distributions: Data, Diagnostics, and Dependence." Tech. Rep. RFF DP 11-19-REV, Resources for the Future.

Costello, C. J., M. G. Neubert, S. A. Polasky, and A. R. Solow. 2010. "Bounded Uncertainty and Climate Change Economics." *Proceedings of the National Academy of Sciences of the United States of America* 107 (18):8108–8110.

Craig, P. P., A. Gadgil, and J. G. Koomey. 2002. "What Can History Teach Us? A Retrospective Examination of Long-Term Energy Forecasts for the United States." *Annual Review of Energy and the Environment* 27:83–118.

Crutzen, P. J. and M. Oppenheimer. 2008. "Learning about Ozone Depletion." *Climatic Change* 89 (1–2):143–154.

Curry, T. E., S. Ansolabehere, and H. Herzog. 2007. "A Survey of Public Attitudes towards Climate Change and Climate Change Mitigation Technologies in the United States: Analyses of 2006 Results." Tech. Rep. MIT LFEE 2007-01 WP, Massachusetts Institute of Technology Laboratory for Energy and the Environment.

Curtright, A. E., M. G. Morgan, and D. W. Keith. 2008. "Expert Assessments of Future Photovoltaic Technologies." *Environmental Science and Technology* 42 (24):9031–9038.

Dantzig, G. B. 1955. "Linear Programming under Uncertainty." *Management Science* 1 (3–4):197–206.

Dapkus, W. D. and T. R. Bowe. 1984. "Planning for New Electric Generation Technologies: A Stochastic Dynamic Programming Approach." *IEEE Transactions on Power Apparatus and Systems* 103 (6):1447–1453.

de Best-Waldhober, M., D. Daamen, and A. Faaij. 2009. "Informed and Uninformed Public Opinions on CO_2 Capture and Storage Technologies in the Netherlands." *International Journal of Greenhouse Gas Control* 3 (3):322–332.

De Cian, E. and M. Tavoni. 2012. "Mitigation Portfolio and Policy Instruments When Hedging against Climate Policy and Technology Uncertainty." *Environmental Modeling and Assessment* 17 (1–2):123–136.

Dessai, S. and M. Hulme. 2004. "Does Climate Adaptation Policy Need Probabilities?" *Climate Policy* 4 (2):107–128.

Dietz, S. 2011. "High Impact, Low Probability? An Empirical Analysis of Risk in the Economics of Climate Change." *Climatic Change* 108 (3):519–541.

Dixit, A. K. and R. S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

DOE/EIA. 1986. "An Analysis of Nuclear Power Plant Construction Costs." Tech. Rep. DOE/EIA-0485, US Department of Energy, Energy Information Administration.

———. 2007. "Annual Energy Outlook 2007." Tech. Rep. DOE/EIA-0383(2007), US Department of Energy, Energy Information Administration.

———. 2009. "Energy Market and Economic Impacts of H.R. 2454, the American Clean Energy and Security Act of 2009." Tech. Rep. SR-OIAF/2009-05, US Department of Energy, Energy Information Administration.

———. 2010. "Updated Capital Cost Estimates for Electricity Generation Plants." US Department of Energy, Energy Information Administration.

———. 2011a. "Annual Energy Outlook 2011." Tech. Rep. DOE/EIA-0383(2011), US Department of Energy, Energy Information Administration.

———. 2011b. "Annual Energy Review 2011." Tech. Rep. DOE/EIA-0384(2011), US Department of Energy, Energy Information Administration.

———. 2011c. "Form EIA-860 Annual Electric Generator Report." US Department of Energy, Energy Information Administration. URL `http://205.254.135.7s/cneaf/electricity/page/eia860.html`.

———. 2012. "Annual Energy Outlook 2012." Tech. Rep. DOE/EIA-0383(2012), US Department of Energy, Energy Information Administration.

DOE/NETL. 2007. "Tracking New Coal-Fired Power Plants: Coal's Resurgence in Electric Power Generation." US Department of Energy, National Energy Technology Laboratory.

———. 2011. "Life Cycle Greenhouse Gas Inventory of Natural Gas Extraction, Delivery, and Electricity Production." Tech. Rep. DOE/NETL-2011/1522, US Department of Energy, National Energy Technology Laboratory.

Dooley, J. J., R. T. Dahowski, C. L. Davidson, S. Bachu, N. Gupta, and J. Gale. 2004. "A CO_2-Storage Supply Curve for North America and Its Implications for the Deployment of Carbon Dioxide Capture and Storage Systems." In *Proceedings of the 7th International Conference on Greenhouse Gas Control Technologies*.

Douglas, M. and A. B. Wildavsky. 1982. *Risk and Culture: An Essay on the Selection of Technical and Environmental Dangers*. Berkeley, CA: University of California Press.

Dowlatabadi, H. and M. G. Morgan. 1993. "Integrated Assessment of Climate Change." *Science* 259 (5103):1813–1932.

Dupačová, J., J. Hurt, and J. Štěpán. 2002. *Stochastic Modeling in Economics and Finance*. Dordrecht, Netherlands: Kulwer Academic Publishers.

Dupačová, J. and K. Sladký. 2002. "Comparison of Multistage Stochastic Programs with Recourse and Stochastic Dynamic Programs with Discrete Time." *ZAMM Journal of Applied Mathematics and Mechanics* 82 (11–12):753–765.

Durand-Lasserve, O., A. Pierru, and Y. Smeers. 2010. "Uncertain Long-Run Emissions Targets, CO_2 Price, and Global Energy Transition: A General Equilibrium Approach." *Energy Policy* 38 (9):5108–5122.

Einsiedel, E. F., A. D. Boyd, J. Medlock, and P. Ashworth. 2013. "Assessing Socio-Technical Mindsets: Public Deliberations on Carbon Capture and Storage in the Context of Energy Sources and Climate Change." *Energy Policy* 53:149–158.

Enders, W. 2009. *Applied Econometric Time Series*. Hoboken, NJ: John Wiley & Sons, Inc.

EPA. 2011a. "Inventory of US Greenhouse Gas Emissions and Sinks: 1990–2009." Tech. Rep. EPA 430-R-11-005, US Environmental Protection Agency.

———. 2011b. "Inventory of US Greenhouse Gas Emissions and Sinks: 1990–2011." Tech. Rep. EPA 430-R-13-001, US Environmental Protection Agency.

Epstein, L. G. 1980. "Decision Making and the Temporal Resolution of Uncertainty." *International Economic Review* 21 (2):269–282.

Evans, D. J. and H. Sezer. 2004. "Social Discount Rates for Six Major Countries." *Applied Economics Letters* 11 (9):557–560.

———. 2005. "Social Discount Rates for Member Countries of the European Union." *Journal of Economic Studies* 32 (1):47–59.

Fairbank, J. K. 2006. *China: A New History.* Cambridge, MA: Harvard University Press.

Fiorese, G., M. Catenacci, E. Verdolini, and V. Bosetti. 2013. "Advanced Biofuels: Future Perspectives from an Expert Elicitation Survey." *Energy Policy* 56:293–311.

Fischer, C. and R. G. Newell. 2008. "Environmental and Technology Policies for Climate Mitigation." *Journal of Environmental Economics and Management* 55 (2):142–162.

Fischer, C., I. W. H. Parry, and W. A. Pizer. 2003. "Instrument Choice for Environmental Protection When Technological Innovation Is Endogenous." *Journal of Environmental Economics and Management* 45 (3):523–545.

Fischhoff, B. 1982. "For Those Condemned to Study the Past: Heuristics and Biases in Hindsight." In *Judgment under Uncertainty: Heuristics and Biases*, edited by D. Kahneman, P. Slovic, and A. Tversky. Cambridge, UK: Cambridge University Press.

Fischhoff, B. and R. Beyth. 1975. "'I Knew It Would Happen:' Remembered Probabilities of Once-Future Things." *Organizational Behavior and Human Performance* 13 (1):1–16.

Fischhoff, B., P. Slovic, and S. Lichtenstein. 1977. "Knowing with Certainty: The Appropriateness of Extreme Confidence." *Journal of Experimental Psychology: Human Perception and Performance* 3 (4):552–564.

Fleishman, L. A., W. Bruine de Bruin, and M. G. Morgan. 2010. "Informed Public Preferences for Electricity Portfolios with CCS and Other Low-Carbon Technologies." *Risk Analysis* 30 (9):1399–1410.

Frame, D. and M. R. Allen. 2008. "Climate Change and Global Risk." In *Global Catastrophic Risks*, edited by N. Bostrom and M. M. Ćirković. Oxford, UK: Oxford University Press.

Gathwaite, P. H., J. B. Kadane, and O'Hagan A. 2005. "Statistical Methods for Eliciting Probability Distributions." *Journal of the American Statistical Association* 100 (470):680–700.

Gerst, M. D., R. B. Howarth, and M. E. Borsuk. 2010. "Accounting for the Risk of Extreme Outcomes in an Integrated Assessment of Climate Change." *Energy* 38 (8):4540–4548.

Global CCS Institute. 2012. "The Global Status of CCS 2012." Global CCS Institute.

Glynn, P. W. and G. Infanger. 2013. "Simulation-Based Confidence Bounds for Two-Stage Stochastic Programs." *Mathematical Programming* 138 (1–2):15–42.

Goolsbee, A. 1998. "Does Government R&D Policy Mainly Benefit Scientists and Engineers?" *American Economic Review* 88 (2):298–302.

Griliches, Z. 1992. "The Search for R&D Spillovers." *Scandinavian Journal of Economics* 94:S29–S47.

Grübler, A. 2010. "The Costs of the French Nuclear Scale-Up: A Case of Negative Learning by Doing." *Energy Policy* 38 (9):5174–5188.

Hall, J. W., R. J. Lempert, K. Keller, A. Hackbarth, C. Mijere, and D. McInerney. 2012. "Robust Climate Policies under Uncertainty: A Comparison of Robust Decision Making and Info-Gap Methods." *Risk Analysis* 32 (10):1657–1672.

Hallegatte, S., A. Shah, R. Lempert, C. Brown, and S. Gill. 2012. "Investment Decision Making under Deep Uncertainty: Application to Climate Change." Tech. Rep. Policy Research Working Paper 6193, The World Bank.

Hammitt, J. K. and A. I. Shlyakhter. 1999. "The Expected Value of Information and the Probability of Surprise." *Risk Analysis* 19 (1):135–152.

Hannart, A., J.-L. Dufresne, and P. Naveau. 2009. "Why Climate Sensitivity May Not Be So Unpredictable." *Geophysical Research Letters* 36 (16).

Hannart, A., M. Ghil, and J.-L. Dufresne. 2013. "Disconcerting Learning on Climate Sensitivity and the Uncertain Future of Uncertainty." *Climatic Change* 119 (3–4):585–601.

Haurie, A., M. Tavoni, and B. C. C. van der Zwaan. 2012. "Modeling Uncertainty and the Economics of Climate Change: Recommendations for Robust Energy Policy." *Environmental Modeling and Assessment* 17 (1–2):1–5.

Heal, G. and A. Millner. 2013. "Uncertainty and Decision in Climate Change Economics." Tech. Rep. 18929, National Bureau of Economic Research.

Heinrich, G., M. Howells, L. Basson, and J. Petrie. 2007. "Electricity Supply Industry Modelling for Multiple Objectives under Demand Growth Uncertainty." *Energy* 32 (11):2210–2229.

Henrion, M. and B. Fischhoff. 1986. "Assessing Uncertainty in Physical Constants." *American Journal of Physics* 54 (9):791–798.

Hoffman, R. R., N. R. Shadbolt, A. M. Burton, and G. Klein. 1995. "Eliciting Knowledge from Experts: A Methodological Analysis." *Organizational Behavior and Human Decision Processes* 62 (2):129–158.

Hogarth, R. M. 1975. "Cognitive Processes and the Assessment of Subjective Probability Distributions." *Journal of the American Statistical Association* 70 (350):271–289.

Hope, C. 2006. "The Marginal Impact of CO_2 from PAGE2002: An Integrated Assessment Model Incorporating the IPCC's Five Reasons for Concern." *Integrated Assessment* 6 (1):19–56.

Hora, S. C. and D. von Winterfeldt. 1997. "Nuclear Waste and Future Societies: A Look into the Deep Future." *Technological Forecasting and Social Change* 56 (2):155–170.

Howard, R. A. 1968. "The Foundations of Decision Analysis." *IEEE Transactions on Systems Science and Cybernetics* 4 (3):211–219.

———. 1971. "Proximal Decision Analysis." *Management Science* 17 (9):507–541.

Howarth, R. W., R. Santoro, and A. Ingraffea. 2011. "Methane and the Greenhouse-Gas Footprint of Natural Gas from Shale Formations." *Climatic Change* 106 (4):679–690.

———. 2012. "Venting and Leaking of Methane from Shale Gas Development: Responses to Cathles et al." *Climatic Change* 113 (2):537–549.

Hu, M.-C. and B. F. Hobbs. 2010. "Analysis of Multi-Pollutant Policies for the US Power Sector under Technology and Policy Uncertainty Using MARKAL." *Energy* 35 (12):5430–5442.

Hultman, N., D. Rebois, M. Scholten, and C. Ramig. 2011. "The Greenhouse Impact of Unconventional Gas for Electricity Generation." *Environmental Research Letters* 6 (4).

Hultman, N. E., J. G. Koomey, and D. M. Kammen. 2007. "What History Can Teach Us about the Future Costs of US Nuclear Power." *Environmental Science and Technology* 41 (7):2087–2094.

Huntington, H. 2013. "EMF 26: Changing the Game?: Emissions and Market Implications of New Natural Gas Supplies." Tech. Rep. EMF Report 26, Volume I, Energy Modeling Forum.

Hwang, I. C., F. Reynès, and R. S. J. Tol. 2011. "Climate Policy under Fat-Tailed Risk: An Application of DICE." Tech. Rep. 403, Economic and Social Research Institute.

IAEA. 2013. "Power Reactor Information System." International Atomic Energy Agency.

IEA. 2012a. "Golden Rules for a Golden Age of Gas." International Energy Agency.

———. 2012b. "World Energy Outlook 2012." International Energy Agency.

IHS CERA. 2013. "Power Capital Cost Index." IHS Cambridge Energy Research Associates. URL `http://www.ihs.com/info/cera/ihsindexes/index.aspx`.

Infanger, G. 1994. *Planning under Uncertainty: Solving Large-Scale Stochastic Linear Programs*. Penryn, CA: Personal Transformation Press.

———. 1999. *GAMS/DECIS User's Guide*.

IPCC. 2000. "Special Report on Emissions Scenarios." Intergovernmental Panel on Climate Change.

———. 2005. "Special Report on Carbon Capture and Storage." Intergovernmental Panel on Climate Change.

———. 2007. "Contribution of Working Groups I, II, and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change." Intergovernmental Panel on Climate Change.

———. 2013. "Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change: The Physical Science Basis." Intergovernmental Panel on Climate Change.

Ishii, J. and J. Yan. 2004. "Investment under Regulatory Uncertainty: US Electricity Generation Investment since 1996." University of California Energy Institute, Center for the Study of Energy Markets.

Jaffe, A. B. 2012. "Technology Policy and Climate Change." *Climate Change Economics* 3 (4).

Jenni, K. E., E. D. Baker, and G. F. Nemet. 2013. "Expert Elicitations of Energy Penalties for Carbon Capture Technologies." *International Journal of Greenhouse Gas Control* 12:136–145.

Jiang, M., W. M. Griffin, C. Hendrickson, P. Jaramillo, J. VanBriesen, and A. Venkatesh. 2011. "Life Cycle Greenhouse Gas Emissions of Marcellus Shale Gas." *Environmental Research Letters* 6 (3).

Jin, S., S. M. Ryan, J.-P. Watson, and D. L. Woodruff. 2011. "Modeling and Solving a Large-Scale Generation Expansion Planning Problem under Uncertainty." *Energy Systems* 2 (3–4):209–242.

Jung, C., K. Krutilla, and R. Boyd. 1996. "Incentives for Advanced Pollution Abatement Technology at the Industry Level: An Evaluation of Policy Alternatives." *Journal of Environmental Economics and Management* 30 (1):95–111.

Kahan, D. M. 2012. "Cultural Cognition as a Conception of the Cultural Theory of Risk." In *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk*, edited by S. Roeser, R. Hillerbrand, P. Sandin, and M. Peterson. London, UK: Springer.

Kahan, D. M., E. Peters, M. Wittlin, P. Slovic, D. Ouellette, L. L. Braman, and G. Mandel. 2012. "The Polarizing Impact of Science Literacy and Numeracy on Perceived Climate Change Risks." *Nature Climate Change* 2:732–735.

Kahneman, D. and A. Tversky. 1982. "Intuitive Prediction: Biases and Corrective Procedures." In *Judgment under Uncertainty: Heuristics and Biases*, edited by D. Kahneman, P. Slovic, and A. Tversky. Cambridge, UK: Cambridge University Press.

———. 2000. "Choices, Values, and Frames." In *Choices, Values, and Frames*, edited by D. Kahneman and A. Tversky. Cambridge, UK: Cambridge University Press.

Kann, A. and J. P. Weyant. 2000. "Approaches to Performing Uncertainty Analysis in Large-Scale Energy/Environmental Policy Models." *Environmental Modeling and Assessment* 5 (1):29–46.

Kanudia, A. and R. Loulou. 1998. "Robust Responses to Climate Change via Stochastic MARKAL: The Case of Québec." *European Journal of Operational Research* 106 (1):15–30.

———. 1999. "Advanced Bottom-Up Modelling for National and Regional Energy Planning in Response to Climate Change." *International Journal of Environment and Pollution* 12 (2–3):191–216.

Kanudia, A. and P. R. Shukla. 1998. "Modelling of Uncertainties and Price Elastic Demands in Energy-Environment Planning for India." *Omega* 26 (3):409–423.

Kaplow, L., E. Moyer, and D. A. Weisbach. 2010. "The Social Evaluation of Intergenerational Policies and Its Application to Integrated Assessment Models of Climate Change." *The B. E. Journal of Economic Analysis and Policy* 10 (2).

Karion, A., C. Sweeney, G. Pétron, G. Frost, R. M. Hardesty, J. Kofler, B. R. Miller, T. Newberger, S. Wolter, R. Banta, A. Brewer, E. Dlugokencky, P. Lang, S. A. Montzka, R. Schnell, P. Tans, M. Trainer, R. Zamora, and S. Conley. 2013. "Methane Emissions Estimate from Airborne Measurements over a Western United States Natural Gas Field." *Geophysical Research Letters* 40 (16):4393–4397.

Katz, R. W. 2002. "Techniques for Estimating Uncertainty in Climate Change Scenarios and Impact Studies." *Climate Research* 20:167–185.

Keats, J. 1899. *The Complete Poetical Works and Letters of John Keats: Cambridge Edition*. Boston, MA: Houghton Mifflin Company.

Keefer, D. L. 1994. "Certainty Equivalents for Three-Point Discrete-Distribution Approximations." *Management Science* 40 (6):760–773.

Keeney, R. L. and D. von Winterfeldt. 1989. "On the Uses of Expert Judgment on Complex Technical Problems." *IEEE Transactions on Engineering Management* 36 (2):83–86.

———. 1991. "Eliciting Probabilities from Experts in Complex Technical Problems." *IEEE Transactions on Engineering Management* 38 (3):191–201.

Keller, K., B. M. Bolker, and D. F. Bradford. 2004. "Uncertain Climate Thresholds and Optimal Economic Growth." *Journal of Environmental Economics and Management* 48 (1):723–741.

Keller, K. and D. McInerney. 2008. "The Dynamics of Learning about a Climate Threshold." *Climate Dynamics* 30 (2–3):321–332.

Kelly, D. L. and C. D. Kolstad. 1999. "Bayesian Learning, Growth, and Pollution." *Journal of Economic Dynamics and Control* 23 (4):491–518.

Kelly, D. L. and Z. Tan. 2013. "Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails." University of Miami.

Keppo, I. and B. van der Zwaan. 2012. "The Impact of Uncertainty in Climate Targets and CO_2 Storage Availability on Long-Term Emissions Abatement." *Environmental Modeling and Assessment* 17 (1–2):177–191.

Kitcher, P. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford, UK: Oxford University Press.

Knight, F. H. 1921. *Risk, Uncertainty, and Profit*. Boston, MA: Houghton Mifflin Company.

Knopf, B., O. Edenhofer, C. Flachsland, M. T. J. Kok, H. Lotze-Campen, G. Luderer, A. Popp, and D. P. van Vuuren. 2010. "Managing the Low-Carbon Transition: From Model Results to Policies." *The Energy Journal* 31 (Special Issue 1: The Economics of Low Stabilization):223–245.

Kolstad, C. D. 1996. "Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions." *Journal of Environmental Economics and Management* 31 (1):1–18.

Koomey, J. G. and N. E. Hultman. 2007. "A Reactor-Level Analysis of Busbar Costs for US Nuclear Plants, 1970–2005." *Energy Policy* 35 (11):5630–5642.

Kousky, C., O. Rostapshova, M. Toman, and R. Zeckhauser. 2010. "Responding to Threats of Climate Change Mega-Catastrophes." Tech. Rep. RWP10-008, Harvard Kennedy School Faculty Research Working Paper Series.

Kriesky, J., B. D. Goldstein, K. Zell, and S. Beach. 2013. "Differing Opinions about Natural Gas Drilling in Two Adjacent Counties with Different Levels of Drilling Activity." *Energy Policy* 58:228–236.

Krukanont, P. and T. Tezuka. 2007. "Implications of Capacity Expansion under Uncertainty and Value of Information: The Near-Term Energy Planning of Japan." *Energy* 32 (10):1809–1824.

Kuhn, T. S. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

Kuhnert, P. M., T. G. Martin, and S. P. Griffiths. 2010. "A Guide to Eliciting and Using Expert Knowledge in Bayesian Ecological Models." *Ecology Letters* 13 (7):900–914.

Kunreuther, H., G. Heal, M. Allen, O. Edenhofer, C. B. Field, and G. Yohe. 2012. "Risk Management and Climate Change." Tech. Rep. 18607, National Bureau of Economic Research.

La Grandville, O. 2012. "How Much Should a Nation Save? A New Answer." *Studies in Nonlinear Dynamics and Econometrics* 16 (2).

Labriet, M., R. Loulou, and A. Kanudia. 2010. "Modeling Uncertainty in a Large Scale Integrated Energy-Climate Model." In *Uncertainty and Environmental Decision Making: A Handbook of Research and Best Practice*, edited by J. A. Filar and A. Haurie. New York, NY: Springer.

Lashof, D. A., S. Yeh, D. Doninger, S. Carter, and L. Johnson. 2012. "Closing the Power Plant Carbon Pollution Loophole: Smart Ways the Clean Air Act Can Clean Up America's Biggest Climate Polluters." Natural Resources Defense Council.

Le Treut, H., R. Somerville, U. Cubasch, Y. Ding, C. Mauritzen, A. Mokssit, T. Peterson, and M. Prather. 2007. "Historical Overview of Climate Change." In *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor, and H. L. Miller. Cambridge, UK: Cambridge University Press.

Leach, A. J. 2007. "The Climate Change Learning Curve." *Journal of Economic Dynamics and Control* 31 (5):1728–1752.

Lempert, R. 2012. "Scenarios That Illuminate Vulnerabilities and Robust Responses." *Climatic Change* 117 (4):627–646.

Lempert, R. and S. McKay. 2011. "Some Thoughts on the Role of Robust Control Theory in Climate-Related Decision Support." *Climatic Change* 107 (3–4):241–246.

Lempert, R. J., D. G. Groves, and J. R. Fischbach. 2013. "Is It Ethical to Use a Single Probability Density Function?" Tech. Rep. WR-992-NSF, RAND Corporation.

Lempert, R. J., S. W. Popper, and S. C. Bankes. 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. Santa Monica, CA: RAND Corporation.

Levi, I. 1974. "On Indeterminate Probabilities." *Journal of Philosophy* 71 (13):391–418.

Levi, M. A. 2012. "Comment on 'Hydrocarbon Emissions Characterization in the Colorado Front Range: A Pilot Study' by Gabrielle Pétron et al." *Journal of Geophysical Research* 117 (D21).

Loulou, R. and A. Kanudia. 1999. "Minimax Regret Strategies for Greenhouse Gas Abatement: Methodology and Application." *Operations Research Letters* 25 (5):219–230.

Loulou, R., M. Labriet, and A. Kanudia. 2009. “Deterministic and Stochastic Analysis of Alternative Climate Targets under Differentiated Cooperation Regimes.” *Energy Economics* 31 (Supplement 2):S131–S143.

Loulou, R. and A. Lehtila. 2012. “Stochastic Programming and Tradeoff Analysis in TIMES.” Energy Technology Systems Analysis Programme.

Luce, R. D. and H. Raiffa. 1957. *Games and Decisions: Introduction and Critical Survey*. New York, NY: John Wiley & Sons, Inc.

Mandansky, A. 1960. “Inequalities for Stochastic Linear Programming Problems.” *Management Science* 6 (2):197–204.

Mandelbrot, B. 1963. “The Variation of Certain Speculative Prices.” *Journal of Business* 36 (4):394–419.

Manne, A. S. 1996. “Hedging Strategies for Global Carbon Dioxide Abatement: A Summary of Poll Results.” Tech. Rep. WP 14.2, Energy Modeling Forum 14 Subgroup (Analysis for Decisions under Uncertainty).

Manne, A. S. and R. G. Richels. 1993. *Buying Greenhouse Insurance: The Economic Costs of CO_2 Emission Limits*. Cambridge, MA: The MIT Press.

Mansfield, E., J. Rapoport, A. Romeo, S. Wagner, and G. Beardsley. 1977. “Social and Private Rates of Return from Industrial Innovations.” *The Quarterly Journal of Economics* 91 (2):221–240.

March, J. G. 1994. *A Primer on Decision Making: How Decisions Happen*. New York, NY: Free Press.

Markowitz, H. 1952. “Portfolio Selection.” *Journal of Finance* 7 (1):77–91.

Martino, J. P. 1987. “Using Precursors as Leading Indicators of Technological Change.” *Technological Forecasting and Social Change* 32 (4):341–360.

Maslin, M. and P. Austin. 2012. “Uncertainty: Climate Models at Their Limit?” *Nature* 486:183–184.

Mason, R. O. 1969. "A Dialectical Approach to Strategic Planning." *Management Science* 15 (8):B403–B414.

Massé, P. and R. Gibrat. 1957. "Application of Linear Programming to Investments in the Electric Power Industry." *Management Science* 3 (2):149–166.

Mastrandrea, M. D. and S. H. Schneider. 2004. "Probabilistic Integrated Assessment of 'Dangerous' Climate Change." *Science* 304 (5670):571–575.

Matheson, D. and J. E. Matheson. 2005. "Describing and Valuing Interventions That Observe or Control Decision Situations." *Decision Analysis* 2 (3):165–181.

Meyer, M. A. and J. M. Booker. 2001. *Eliciting and Analyzing Expert Judgment: A Practical Guide*. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Milliman, S. R. and R. Price. 1989. "Firm Incentives to Promote Technological Change in Pollution Control." *Journal of Environmental Economics and Management* 17 (3):247–265.

Millner, A. 2013. "On Welfare Frameworks and Catastrophic Climate Risks." *Journal of Environmental Economics and Management* 65 (2):310–325.

Millner, A., Calel R., Stainforth D. A., and G. MacKerron. 2013. "Do Probabilistic Expert Elicitations Capture Scientists' Uncertainty about Climate Change?" *Climatic Change* 116 (2):427–436.

Moniz, E. J., H. D. Jacoby, and A. J. M. Meggs. 2010. "The Future of National Gas: Interim Report." Massachusetts Institute of Technology.

Morgan, M. G. 2011. "Certainty, Uncertainty, and Climate Change." *Climatic Change* 108 (4):707–721.

Morgan, M. G., H. Dowlatabadi, M. Henrion, D. Keith, R. Lempert, S. McBride, M. Small, and T. Wilbanks. 2009. "Best Practice Approaches for Characterizing, Communicating, and Incorporating Scientific Uncertainty in Climate Decision Making." US Climate Change Science Program and the Subcommittee on Global Change Research, National Oceanic and Atmospheric Administration.

Morgan, M. G., B. Fischhoff, A. Bostrom, and C. J. Atman. 2002. *Risk Communication: A Mental Models Approach.* Cambridge, UK: Cambridge University Press.

Morgan, M. G. and M. Henrion. 1990. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge, UK: Cambridge University Press.

Morgan, M. G., M. Kandlikar, J. Risbey, and H. Dowlatabadi. 1999. "Why Conventional Tools for Policy Analysis Are Often Inadequate for Problems of Global Change." *Climatic Change* 41 (3–4):271–281.

Morgan, M. G. and D. W. Keith. 1995. "Subjective Judgments by Climate Experts." *Environmental Science and Technology* 29 (10):468A–476A.

Munson, R. 2005. *From Edison to Enron: The Business of Power and What It Means for the Future of Electricity.* Westport, CT: Praeger.

National Research Council. 2005. "Prospective Evaluation of Applied Energy Research and Development at DOE (Phase One): A First Look Forward." The National Academies Press.

———. 2007. "Prospective Evaluation of Applied Energy Research and Development at DOE (Phase Two)." The National Academies Press.

NEA. 2012. "Nuclear Energy Today: Second Edition." Tech. Rep. NEA No. 6885, Organisation for Economic Co-Operation and Development, Nuclear Energy Agency.

Nemet, G. F. 2006. "Beyond the Learning Curve: Factors Influencing Cost Reductions in Photovoltaics." *Energy Policy* 34 (17):3218–3232.

Niemeyer, V., G. Blanford, F. de la Chesnaye, J. Merrick, and D. Young. 2012. "Quantifying the Economic Costs of Electric Sector Environmental Controls." In *31st USAEE/IAEE North American Conference.* Austin, TX.

Nordhaus, W. D. 1994. *Managing the Global Commons: The Economics of Climate Change.* Cambridge, MA: The MIT Press.

———. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press.

———. 2009. "An Analysis of the Dismal Theorem." Tech. Rep. Cowles Foundation Discussion Paper No. 1686, Yale University, Cowles Foundation for Research in Economics.

———. 2012. "Economic Policy in the Face of Severe Tail Events." *Journal of Public Economic Theory* 14 (2):197–219.

NRC. 2010. "Reactor Safety Study: An Assessment of Accident Risks in US Commercial Nuclear Power Plants." US Nuclear Regulatory Commission.

O'Hagan, A. 1998. "Eliciting Expert Beliefs in Substantial Practical Applications." *Journal of the Royal Statistical Society* 47 (1):21–35.

O'Hagan, A., C. E. Buck, A. Daneshkhah, J. R. Eiser, P. H. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow. 2006. *Uncertain Judgments: Eliciting Experts' Probabilities*. West Sussex, UK: John Wiley & Sons, Inc.

OMB. 2003. "Circular A-4." Office of Management and Budget. URL `http://http://www.whitehouse.gov/omb/circulars_a004_a-4`.

O'Neill, B. C., P. Crutzen, A. Grübler, M. H. Duong, K. Keller, C. Kolstad, J. Koomey, A. Lange, M. Obersteiner, M. Oppenheimer, W. Pepper, W. Sanderson, M. Schlesinger, N. Treich, A. Ulph, M. Webster, and C. Wilson. 2006. "Learning and Climate Change." *Climate Policy* 6 (5):585–589.

Oppenheimer, M., B. C. O'Neill, and M. Webster. 2008. "Negative Learning." *Climatic Change* 89 (1–2):155–172.

Oreskes, N. and E. M. Conway. 2010. *Merchants of Doubt*. New York, NY: Bloomsbury Press.

OTA. 1984. "Nuclear Power in an Age of Uncertainty." Tech. Rep. OTA-E-216, US Congress, Office of Technology Assessment.

Pacala, S. and R. Socolow. 2004. "Stabilization Wedges: Solving the Climate Problem for the Next 50 Years with Current Technologies." *Science* 305 (5686):968–972.

Palmgren, C. R., M. G. Morgan, W. Bruine de Bruin, and D. W. Keith. 2004. "Initial Public Perceptions of Deep Geological and Oceanic Disposal of Carbon Dioxide." *Environmental Science and Technology* 38 (24):6441–6450.

Parson, E. A. 2003. *Protecting the Ozone Layer: Science and Strategy*. Oxford, UK: Oxford University Press.

———. 2007. "The Big One: A Review of Richard Posner's 'Catastrophe: Risk and Response'." *Journal of Economic Literature* 45 (1):147–164.

Parson, E. A., V. Burkett, K. Fisher-Vanden, D. Keith, L. Mearns, H. Pitcher, C. Rosenzweig, and M. Webster. 2007. "Global-Change Scenarios: Their Development and Use." Tech. Rep. Sub-Report 2.1B, US Climate Change Science Program and the Subcommittee on Global Change Research, Department of Energy.

Parson, E. A. and D. Karwat. 2011. "Sequential Climate Change Policy." *Wiley Interdisciplinary Reviews: Climate Change* 2 (5):744–756.

Patiño Echeverri, D., P. Fischbeck, and E. Kriegler. 2009. "Economic and Environmental Costs of Regulatory Uncertainty for Coal-Fired Power Plants." *Environmental Science and Technology* 43 (3):578–584.

Peck, S. C. and T. J. Teisberg. 1993. "Global Warming Uncertainties and the Value of Information: An Analysis Using CETA." *Resource and Energy Economics* 15 (1):71–97.

Pétron, G., G. Frost, B. R. Miller, A. I. Hirsch, S. A. Montzka, A. Karion, M. K. Trainer, C. Sweeney, A. E. Andrews, L. Miller, J. Kofler, A. Bar-Ilan, E. J. Dlugokencky, L. Patrick, C. T. Moore, T. B. Ryerson, C. Siso, W. Kolodzey, P. M. Lang, T. J. Conway, P. Novelli, K. Masarie, B. Hall, D. Guenther, D. Kitzis, J. Miller, D. Welsh, D. Wolfe, W. Neff, and P. Tans. 2012. "Hydrocarbon Emissions Characterization in the Colorado Front Range: A Pilot Study." *Journal of Geophysical Research* 117 (D4).

Pétron, G., G. J. Frost, M. K. Trainer, B. R. Miller, E. J. Dlugokencky, and P. Tans. 2013. "Reply to Comment on 'Hydrocarbon Emissions Characterization in the Colorado Front Range: A Pilot Study' by Michael A. Levi." *Journal of Geophysical Research* 118 (1):236–242.

Phillips, L. D. 1999. "Group Elicitation of Probability Distributions: Are Many Heads Better Than One?" In *Decision Science and Technology: Reflections on the Contributions of Ward Edwards*, edited by J. Shanteau, B. Mellors, and D. Schum. Norwell, MA: Kluwer Academic Publishers.

Pindyck, R. S. 2012. "Uncertain Outcomes and Climate Change Policy." *Journal of Environmental Economics and Management* 63 (3):289–303.

Posner, R. A. 2004. *Catastrophe: Risk and Response*. Oxford, UK: Oxford University Press.

Postma, T. J. B. M. and F. Liebl. 2005. "How to Improve Scenario Analysis as a Strategic Management Tool?" *Technological Forecasting and Social Change* 72 (2):161–173.

Pycroft, J., L. Vergano, C. Hope, D. Paci, and J. C. Ciscar. 2011. "A Tale of Tails: Uncertainty and the Social Cost of Carbon Dioxide." *Economics* 5.

Rai, V., D. G. Victor, and M. C. Thurber. 2010. "Carbon Capture and Storage at Scale: Lessons from the Growth of Analogous Energy Technologies." *Energy Policy* 38 (8):4089–4098.

Raiffa, H. and R. Schlaifer. 1961. *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University Press.

Rao, A. B., E. S. Rubin, D. W. Keith, and M. G. Morgan. 2006. "Evaluation of Potential Cost Reductions from Improved Amine-Based CO_2 Capture Systems." *Energy Policy* 34 (18):3765–3772.

Reichelstein, S. and M. Yorston. 2013. "The Prospects for Cost Competitive Solar PV Power." *Energy Policy* 55:117–127.

Reiner, D., T. Curry, M. de Figueiredo, H. Herzog, S. Ansolabehere, K. Itaoka, M. Akai, F. Johnsson, and M. Odenberger. 2006. "An International Comparison of Public Attitudes Towards Carbon Capture and Storage Technologies." In *8th International Conference on Greenhouse Gas Control Technologies.* Trondheim, Norway.

Roe, G. H. and M. B. Baker. 2007. "Why Is Climate Sensitivity So Unpredictable?" *Science* 318 (5850):629–632.

———. 2011. "Comment on 'Another Look at Climate Sensitivity' by Zaliapin and Ghil (2010)." *Nonlinear Processes in Geophysics* 18:125–127.

Roughgarden, T. and S. H. Schneider. 1999. "Climate Change Policy: Quantifying Uncertainties for Damages and Optimal Carbon Taxes." *Energy Policy* 27 (7):415–429.

Rowlands, D. J., D. J. Frame, D. Ackerley, T. Aina, B. B. B. Booth, C. Christensen, M. Collins, N. Faull, C. E. Forest, B. S. Grandley, E. Gryspeerdt, E. J. Highwood, W. J. Ingram, S. Knight, A. Lopez, N. Massey, F. McNamera, N. Meinshausen, C. Piani, S. M. Rosier, B. M. Sanderson, L. A. Smith, D. A. Stone, M. Thurston, K. Yamazaki, Y. H. Yamazaki, and M. R. Allen. 2012. "Broad Range of 2050 Warming from an Observationally Constrained Large Climate Model Ensemble." *Nature Geoscience* 5:256–260.

Rubin, E. S., S. Yeh, M. Antes, M. Berkenpas, and J. Davison. 2007. "Use of Experience Curves to Estimate the Future Cost of Power Plants with CO_2 Capture." *International Journal of Greenhouse Gas Control* 1 (2):188–197.

Ryan, S. M., J. D. McCalley, and D. L. Woodruff. 2011. "Long Term Resource Planning for Electric Power Systems under Uncertainty." Lawrence Berkeley National Laboratory for Computation of the Next Generation Grids.

Savage, L. J. 1954. *The Foundations of Statistics.* Hoboken, NJ: John Wiley & Sons, Inc.

Savage, S. 2009. *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty.* Hoboken, NJ: John Wiley & Sons, Inc.

Schneider, S. H. and K. Kuntz-Duriseti. 2002. "Uncertainty and Climate Change Policy." In *Climate Change Policy: A Survey*, edited by S. H. Schneider, A. Rosencranz, and J. O. Niles. Washington, DC: Island Press.

Seaver, D. A. 1978. "Assessing Probability with Multiple Individuals: Group Interaction Versus Mathematical Aggregation." Tech. Rep. Report No. 78-3, Social Science Research Institute.

Shapiro, A., D. Dentcheva, and A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory.* Philadelphia, PA: Society for Industrial and Applied Mathematics and the Mathematical Programming Society.

Shindell, D. T., G. Faluvegi, D. M. Kock, G. A. Schmidt, N. Unger, and S. E. Bauer. 2009. "Improved Attribution of Climate Forcing to Emissions." *Science* 326 (5953):716–718.

Shlyakhtera, A. I., D. M. Kammen, C. L. Broido, and R. Wilson. 1994. "Quantifying the Credibility of Energy Projections from Trends in Past Data: The US Energy Sector." *Energy Policy* 22 (2):119–130.

Small, M. J. and P. S. Fischbeck. 1999. "False Precision in Bayesian Updating with Incomplete Models." *Human and Ecological Risk Assessment: An International Journal* 5 (2):291–304.

Söderholm, P. and T. Sundqvist. 2007. "Empirical Challenges in the Use of Learning Curves for Assessing the Economic Prospects of Renewable Energy Technologies." *Renewable Energy* 32 (15):2559–2578.

Solomon, S., G.-K. Plattner, R. Knutti, and P. Friedlingstein. 2009. "Irreversible Climate Change Due to Carbon Dioxide Emissions." *Proceedings of the National Academy of Sciences of the United States of America* 106 (6):1704–1709.

Staël von Holstein, C.-A. S. and J. E. Matheson. 1979. *A Manual for Encoding Probability Distributions*. Palo Alto, CA: SRI International.

Stainforth, D. A., T. Aina, C. Christensen, M. Collins, N. Faull, D. J. Frame, J. A. Kettleborough, S. Knight, A. Martin, J. M. Murphy, C. Piani, D. Sexton, L. A. Smith, R. A. Spicer, A. J. Thorpe, and M. R. Allen. 2005. "Uncertainty in Predictions of the Climate Response to Rising Levels of Greenhouse Gases." *Nature* 433:403–406.

Stern, N. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.

Sunstein, C. R. 2007. *Worst-Case Scenarios*. Cambridge, MA: Harvard University Press.

Taleb, N. N. 2007. *The Black Swan: The Impact of the Highly Improbable*. New York, NY: Random House.

———. 2012. *Antifragile: Things That Gain from Disorder*. New York, NY: Random House.

Tarui, N. and S. Polasky. 2005. "Environmental Regulation with Technology Adoption, Learning, and Strategic Behavior." *Journal of Environmental Economics and Management* 50 (3):447–467.

Taylor, M. R. 2005. "Regulation as the Mother of Innovation: The Case of SO_2 Control." *Law and Policy* 27 (2):348–378.

———. 2012. "Innovation under Cap-and-Trade Programs." *Proceedings of the National Academy of Sciences of the United States of America* 109 (13):4804–4809.

Tol, R. S. J. 2003. "Is the Uncertainty about Climate Change Too Large for Expected Cost-Benefit Analysis?" *Climatic Change* 56 (3):265–289.

Tomassini, L., P. Reichert, R. Knutti, T. F. Stocker, and M. E. Borsuk. 2007. "Robust Bayesian Uncertainty Analysis of Climate System Properties Using Markov Chain Monte Carlo Methods." *Journal of Climate* 20 (7):1239–1254.

Traeger, C. P. 2009. "Recent Developments in the Intertemporal Modeling of Uncertainty." *Annual Review of Resource Economics* 1 (1):261–286.

Tuthill, L. 2008. "Investment in Electricity Generation under Emissions Price Uncertainty: The Plant-Type Decision." Tech. Rep. EV 39, Oxford Institute for Energy Studies.

Tversky, A. and D. Kahneman. 1982. "Judgment under Uncertainty: Heuristics and Biases." In *Judgment under Uncertainty: Heuristics and Biases*, edited by D. Kahneman, P. Slovic, and A. Tversky. Cambridge, UK: Cambridge University Press.

Unger, D. and H. Herzog. 1998. "Comparative Study on Energy R&D Performance: Gas Turbine Case Study." Central Research Institute of Electric Power Industry.

United States House of Representatives. 111th Congress. 1st Session. 2009. "American Clean Energy and Security Act of 2009 (H.R. 2454)." URL `http://www.gpo.gov/fdsys/pkg/BILLS-111hr2454ih/pdf/BILLS-111hr2454ih.pdf`.

Usher, W. and N. Strachan. 2012. "Critical Mid-Term Uncertainties in Long-Term Decarbonisation Pathways." *Energy Policy* 41:433–444.

Van Noorden, R. 2010. "Carbon Sequestration: Buried Trouble." *Nature* 463:871–873.

von Neumann, J. and O. Morgenstern. 1944. *Theory of Games and Economic Behavior.* Princeton, NJ: Princeton University Press.

Wagner, G. and R. J. Zeckhauser. 2012. "Climate Policy: Hard Problem, Soft Thinking." *Climatic Change* 110 (3–4):507–521.

Wallace, S. W. and S.-E. Fleten. 2003. "Stochastic Programming Models in Energy." *Handbooks in Operations Research and Management Science* 10:637–677.

Wallace, S. W. and W. T. Ziemba, editors. 2005. *Applications of Stochastic Programming.* Philadelphia, PA: Society for Industrial and Applied Mathematics and the Mathematical Programming Society.

Webster, M. D. 2002. "The Curious Role of 'Learning' in Climate Policy: Should We Wait for More Data?" *The Energy Journal* 23 (2):97–119.

Webster, M. D., L. Jakobovits, and J. Norton. 2008. "Learning about Climate Change and Implications for Near-Term Policy." *Climatic Change* 89 (1–2):67–85.

Webster, M. D., A. P. Sokolov, J. M. Reilly, C. E. Forest, S. Paltsev, A. Schlosser, C. Wang, D. Kicklighter, M. Sarofim, J. Melillo, R. G. Prinn, and H. D. Jacoby. 2012. "Analysis of Climate Policy Targets under Uncertainty." *Climatic Change* 112 (3–4):569–583.

Weitzman, M. L. 2007. "A Review of *The Stern Review on the Economics of Climate Change*." *Journal of Economic Literature* 45 (3):703–724.

———. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91 (1):1–19.

———. 2012. "GHG Targets as Insurance against Catastrophic Climate Damages." *Journal of Public Economic Theory* 14 (2):221–244.

Weyant, J. P. 2004. "Introduction and Overview: EMF-19 Alternative Technology Strategies for Climate Change Policy." *Energy Economics* 26 (4):501–515.

———. 2008. "A Critique of the *Stern Review's* Mitigation Cost Analyses and Integrated Assessment." *Review of Environmental Economics and Policy* 2 (1):77–93.

Weyant, J. P., F. C. de la Chesnaye, and G. J. Blanford. 2006. "Overview of EMF-21: Multigas Mitigation and Climate Policy." *The Energy Journal* (Multi-Greenhouse Gas Mitigation and Climate Policy Special Issue):1–32.

Wilkerson, J., P. Larsen, and G. Barbose. 2013. "Survey of Western US Electric Utility Resource Plans." *Energy Policy* (Forthcoming).

Yohe, G., N. Andronova, and M. Schlesinger. 2004. "To Hedge or Not against an Uncertain Climate Future?" *Science* 306:416–417.

Yohe, G. W. and R. S. J. Tol. 2007. "Precaution and a Dismal Theorem: Implications for Climate Policy and Climate Research." Tech. Rep. FNU-145, Hamburg University and Centre for Marine and Atmospheric Science.

Zaliapin, I. and M. Ghil. 2010. "Another Look at Climate Sensitivity." *Nonlinear Processes in Geophysics* 17:113–122.

Zdybel, R. and E. Baker. 2013. "Generation of a Correlated Probability Distribution for End-User Energy Prices." *Energy Economics* (Under Revision).

Zickfield, K., M. G. Morgan, D. J. Frame, and D. W. Keith. 2010. "Expert Judgments about Transient Climate Response to Alternative Future Trajectories of Radiative Forcing." *Proceedings of the National Academy of Sciences of the United States of America* 107 (28):12451–12456.