Electricity Generation and Emissions Reduction Decisions under Uncertainty: A General Equilibrium Analysis

by

Jennifer Faye Morris

M.S. Technology and Policy
Massachusetts Institute of Technology, 2009

B.A., Public Policy Analysis and History
University of North Carolina at Chapel Hill, 2006

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2013

Signature of Author………………………………………………………………………………...
Engineering Systems Division
August 30, 2013

Certified by………………………………………………………………...........................
Mort D. Webster
Assistant Professor of Engineering Systems
Committee Chair

Certified by………………………………………………………………...........................
Joseph M. Sussman
JR East Professor of Civil and Environmental Engineering and Engineering Systems
Committee Member

Certified by………………………………………………………………...........................
John M. Reilly
Senior Lecturer, Sloan School of Management
Committee Member

Accepted by………………………………………………………………..........................
Richard C. Larson
Mitsui Professor of Engineering Systems
Chair, Engineering Systems Division Education Committee

(This page is intentionally left blank.)

Electricity Generation and Emissions Reduction Decisions under Uncertainty: A General Equilibrium Analysis

by

Jennifer Faye Morris

Submitted to the Engineering Systems Division
on August 31, 2013 in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Engineering Systems

ABSTRACT

The electric power sector, which accounts for approximately 40% of U.S. carbon dioxide emissions, will be a critical component of any policy the U.S. government pursues to confront climate change. In the context of uncertainty in future policy limiting emissions and future technology costs, society faces the following question: What should the electricity mix we build in the next decade look like? We can continue to focus on conventional generation or invest in low-carbon technologies. There is no obvious answer without explicitly considering the risks created by uncertainty.

This research investigates socially optimal near-term electricity investment decisions under uncertainty in future policy and technology costs. It employs a novel framework that models decision-making under uncertainty with learning in an economy-wide setting that can measure social welfare impacts. Specifically, a computable general equilibrium (CGE) model is formulated as a two-stage stochastic dynamic program focused on decisions in the electric power sector.

The new model is applied to investigate a number of factors affecting optimal near-term electricity investments: (1) policy uncertainty, (2) expansion rate limits on low-carbon generation, (3) low-carbon technology cost uncertainty, (4) technological learning (i.e., near-term investment lowers the expected future technology cost), and (5) the inclusion of a safety valve in future policy which allows the emissions cap to be exceeded, but at a cost.

In modeling decision-making under uncertainty, an optimal electricity investment hedging strategy is identified. Given the experimental design, the optimal hedging strategy reduces the expected policy costs by over 50% compared to a strategy derived using the expected value for the uncertain parameter; and by 12-400% compared to strategies developed under a perfect foresight or myopic framework.

This research also shows that uncertainty has a cost, beyond the cost of meeting a policy. In the experimental design used here, uncertainty in the future policy increases the expected cost of policy by over 45%. If political consensus can be reached and the climate science uncertainties resolved, setting clear, long-term policies can minimize expected policy costs.

In addition, this work contributes to the learning-by-doing literature by presenting a stochastic formulation of technological learning in which near-term investments in a technology affect the probability distribution of the future cost of that technology. Results using this formulation demonstrate that learning rates lower than those found in the literature can lead to significant additional near-term investment in low-carbon technology in order to lower the expected future cost of the technology in case a stringent policy is adopted.

Ultimately, this dissertation demonstrates that near-term investments in low-carbon technologies should be greater than what would be justified to meet near-term goals alone. Near-term low-carbon investments can lower the expected cost of future policy by developing a less carbon-intensive electricity mix, spreading the burden of emissions reductions over time, helping to overcome technology expansion rate constraints, and reducing the expected future cost of low-carbon technologies—all of which provide future flexibility in meeting a policy. The additional near-term cost of low-carbon investments is justified by the future flexibility that such investments create. The value of this flexibility is only explicitly considered in the context of decision-making under uncertainty.

Thesis Supervisor: Mort D. Webster
Assistant Professor of Engineering Systems

Committee Member: Joseph M. Sussman
JR East Professor of Civil and Environmental Engineering and Engineering Systems

Committee Member: John M. Reilly
Senior Lecturer, Sloan School of Management

Acknowledgements

I cannot express enough gratitude to my incredible adviser Mort Webster. I first met Mort as a sophomore in undergrad at UNC, and he has been my mentor ever since. For that, I am truly lucky. He has opened the door to opportunities that I never would have imagined, and has supported and encouraged me throughout. Mort has played a central role in my development as researcher, as well as a person. I am inspired by his passion for his work, his dedication to his students and teaching, and the value he places on a balanced life. I have always been able to turn to him for advice, feedback, motivation, and help, both in the context of my research and life in general. Mort is a great role model, and I look forward to working together in the future.

I also thank Joe Sussman. As a committee member, Joe has been an excellent source of support and ideas. I truly appreciate the way he encourages me to think about alternative possible perspectives, explanations, outcomes, etc. and reminds me of the big picture.

I also owe tremendous thanks to John Reilly, who has been a wonderful adviser throughout my time at MIT. John provides a rare mix of priceless wisdom and humor, which makes working with him very enjoyable. He always offers great insight, compelling hypotheses, and a penchant for precision. I am excited to continue to work with him in the future.

The Joint Program on the Science and Policy of Global Change at MIT has been a truly amazing place to work. I cannot imagine a better group of people to work with and learn from every day. I extend heartfelt appreciation for the invaluable help and support given to me by Sergey Paltsev, who has taught me a great deal. I also need to thank Fannie Barnes and Tony Tran for the fantastic assistance and company they provide to us all at the Joint Program. I further thank the Joint Program sponsors, including the U.S. Department of Energy, the U.S. Environmental Protection Agency, and a consortium of industry and foundation sponsors.

In addition, I would like to thank my JP and ESD friends over the years who have made my time here as a student so enjoyable. I would like to especially acknowledge Valerie Karplus, Nidhi Santen, Bryan Palmintier, Pearl Donohoo, Maite Peña Alcaraz, Claudia Octaviano, Rebecca Saari, DJ Ramberg, Vivek Sakhrani, Fernando de Sisternes, Paul Kishimoto and Chris Gillespie, among many others. Also thanks to Beth Milnes who takes care of us ESD students.

Finally, I would like to thank my family and friends for their support and all of the fun times which gave me a break from work and school. As my time as a student comes to an end, I thank my Mom and Dad for their endless support and encouragement of my academic pursuits. I could not have done it without them. I thank them, my brothers, in-laws, and extended family for their enduring love and belief in me. As well as for rallying around me in the end and helping to take care of Evelyn.

I particularly thank my husband Josh for his amazing love, support, and encouragement, and his patience in listening to me talk about general equilibrium modeling and stochastic dynamic programming. I got through this dissertation because of him. He picked up the slack on household chores (I can't remember the last time I did dishes!), he kept me grounded, he kept me laughing, and he helped me keep balance in my life, which helped me throughout the PhD program. Josh, you are an amazing partner and I love you.

I dedicate this dissertation to my amazing daughter Evelyn who I love so much. Evie provided so much joy, perspective and motivation at the end of the PhD process. I know her curious, tenacious, spirited, and exploratory nature will do her well in an uncertain and complex world. It is my hope that she and those of her generation can make wise decisions for our world and planet.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

1.1 The Problem

As the United States considers its options for reducing greenhouse gas (GHG) emissions to confront climate change, it is clear that the electric power sector will be critical component of any emissions reduction policy. In the U.S. electric power generation is responsible for approximately 40% of all carbon dioxide (CO_2) emissions. To reduce electricity emissions society must switch to cleaner energy sources for generation and/or reduce overall energy use by reducing consumption or increasing efficiency. The government and individual investors must make these capacity expansion and emission reduction decisions today while future policy, technology costs, and other factors remain uncertain. In addition to uncertainty, decisions affecting the evolution of the electricity system are particularly difficult because of the size and long lifetimes of investments. A large coal plant can cost over $4 billion, and a moderate-sized wind farm can cost over $1 billion (EIA, 2013). The lifetimes of plants are at least 40 years, and in some cases over 60 years. As a result new capacity investment decisions made in the near-term are linked to operation decisions for decades to come. In terms of meeting long-term environmental goals, emissions are the net effect of both of these decisions. Before making expensive, long-lived investments that may be subject to future policy, it is wise to consider how to minimize the risks created by uncertainty.

Of the many uncertainties involved in electricity decision-making, uncertainty in future government climate policy is of immense importance as it affects the solvency of long-lived capacity investments. If a climate policy is implemented during the early or mid lifetime of a power plant, it greatly affects how cost-effective it is to run that plant and could even require the plant to shut down, which in turn affects how cost-effective it is to build that plant in the first place. Another important uncertainty is in the costs of technologies, which drive investment decisions and in turn emissions and the cost of reducing emissions. Currently many zero- or low-emitting generation technologies, such as wind, solar, nuclear, and coal or natural gas with carbon capture and storage (CCS), are relatively more expensive than conventional coal and natural gas generation and may not yet be commercially available. However, investment now could go a long way in bringing down the cost for the future, providing greater flexibility and ease in meeting future policy. Hence, investment in these advanced low-carbon technologies and

emissions reduction strategies may only make sense today if we consider the prospect of future policy or see investment in the current technology as a way to lower its cost in the future.

In the face of policy and technology cost uncertainty, society faces these very challenging questions: what should the electricity mix we build in the next decade look like? Should we start building more renewables now? Should we begin to focus on CCS or nuclear? Should we postpone investment in advanced carbon-free technologies until later? Should we rely on more natural gas? There is no obvious answer to these questions without calculating optimal investment under uncertainty. We can wait until policy is implemented before investing in expensive advanced low-carbon technologies and continue to invest in conventional fossil technologies in the meantime. In particular, natural gas, a cleaner burning fuel than coal, is commonly championed as the "bridge" to a less carbon-intensive electricity mix. These "wait and see" strategies are cheaper in the near-term, and would be smart if ultimately no policy were implemented. However, future policy may make these poor strategies. In that case, near-term decisions that create a fossil intensive electricity mix would make it very costly to significantly reduce emissions in the future, potentially even forcing fossil plants to shut down before the end of their lifetime. Even a system dominated by natural gas, a relatively inexpensive but lower carbon option than coal, could make for costly reductions in the face of stringent policy.

On the other hand, as a society we could begin investing in the more expensive zero- or low-emitting technologies today. This would increase near-term costs, but would make future policy cheaper to meet. Current investment in low-carbon technologies would create a less carbon-intensive electricity mix which would help in meeting future policy. Additionally, it could potentially lower the costs of low-carbon technologies in the future, through learning and scale effects. Large shares of low-carbon technologies may be required, depending on the stringency of future policy, and investments now to make those technologies cheaper would make meeting the policy less costly. However, if ultimately there was no policy or a weaker policy, low-carbon technologies may be unnecessary. In that case, near-term investments would be wasted. Investments could also be uneconomical if the cost of the low-carbon technology did not decrease in the future despite near-term investments.

Which of these strategies is best from a social welfare perspective is unclear and depends on intertemporal tradeoffs in terms of costs, which can only be evaluated based on how the policy uncertainty unfolds. The question can be thought of as one of "now vs. later". Investing in

expensive low-carbon technologies now leads to higher near-term costs but lower future costs of meeting a climate policy. That tradeoff is only worth it if the policy will be stringent enough. Investing in cheaper conventional fossil technologies leads to lower near-term costs but higher future costs of meeting a policy. That tradeoff is worth it unless the policy will be stringent. The complexity created by intertemporal tradeoffs and uncertainty in future policy and technology costs makes it difficult to determine an appropriate generation expansion strategy. Essentially we are left to use the imperfect information we have available to place bets on different technologies, which pay off according to how the uncertainties unfold.

The complexity of electricity generation decision making calls for decision support tools. Bottom-up engineering cost models and partial equilibrium models focused on the electricity sector are common tools employed for such decision making. Another powerful type of tool is a computable general equilibrium (CGE) model, which has the advantage of being able to capture all of the interactions and indirect effects throughout an economy and to calculate economy-wide policy costs. While these different decision support models can be quite useful, they tend to be deterministic. Failing to explicitly capture decision-making under uncertainty has been a major drawback of existing decision support models. In fact, the U.S. Department of Energy (DOE) Office of Science has declared that representation of uncertainty and risk is one of the key scientific challenges that must be overcome to further improve CGE and integrated assessment models and their relevance as real-world decision aids (Janetos *et al.*, 2009).

The absence of decision-making under uncertainty in CGE and other models also relates to how the expectations of decision makers are represented. The majority of economic models assume either of two assumptions. Some models assume forward-looking behavior with perfect foresight, in which the future is assumed to be known for certain and decisions over the entire time horizon are optimized at once for a single scenario. Others are myopic recursive-dynamic, in which it assumed that nothing will change in the future and decisions are made in each time period sequentially without any information on what will happen in future periods. Neither of these representations of expectations reflects the real world. Decision-makers have uncertain expectations about the future that are updated as new information is acquired, and decisions can be revised. A stochastic dynamic model reflects this reality by formally factoring these expectations about uncertain quantities into decision-making.

In the real world, public and private investors struggle with how to incorporate uncertainty into decision-making. Investors do their best to make use of the information they have and gather new information. They often pursue activities like carbon price forecasting and emissions profiling. Commonly, uncertainty is addressed through scenario analysis or parametric sensitivity analysis. However, in such analyses each scenario is still deterministic and each will suggest a different optimal investment strategy. In this way, they do not provide actionable information—they say what to do for a particular scenario but not what to do if we do not know which scenario we are in or will be in. A common way to get around this is to take an average of the scenarios or the middle scenario and base decisions on that. However in a complex, nonlinear system like electricity, this approach suffers from the flaw of averages and does not properly capture risks. It is also common to simply ignore uncertainty, perhaps believing that information is too uncertain to make use of. However, there is value in properly utilizing imperfect information and formally considering uncertainty in decision making. While there are sectoral models of the electric power sector that do include decision-making under uncertainty, there is very little done in an economy-wide framework. Doing so can help reframe electricity decisions as risk management and identify socially optimal hedging strategies. Such strategies are particularly important in the electric power sector where generation investment decisions are relatively "irreversible" and have a large impact on long-term environmental goals.

1.2 Purpose of this Dissertation

The purpose of this dissertation is to present a new decision-support framework that addresses the challenges discussed above and can investigate socially optimal capacity investment and emission reduction decisions in the electric power sector under uncertainty. There have been sector-specific studies that capture decision-making under uncertainty well, but cannot address economy-wide social welfare implications. There have also been economy-wide computable general equilibrium (CGE) studies with uncertainty, but without capturing the critical nature of making decisions under uncertainty, learning, and then making decisions again. This work provides a new framework that models decision-making under uncertainty with learning and the ability to revise decisions over time in an economy-wide setting that can measure social welfare impacts. It does this by merging two methodologies: computable general equilibrium (CGE) economic modeling and dynamic programming. By formulating a CGE

model as a two-stage stochastic dynamic program focused on decisions in the electric power sector, this dissertation investigates electricity investment decisions under uncertainty. When uncertainty is explicitly considered, questions of electric generation capacity expansion and emissions reductions become fundamentally questions of risk management and hedging against future costs.

The new framework considers economy-wide effects of electric power sector decisions and policies, intertemporal tradeoffs, and uncertainty in future government policy and technology costs. It also considers the complementary and linked roles that investment in new generation capacity and operation of existing capacity play. Results obtained from the decision-support model can provide insight and information about socially optimal near-term electricity investment strategies that will hedge against the risks associated with uncertainty.

The model framework itself demonstrates how CGE models can capture stochastic dynamic expectations and facilitates further model development and analysis in this area. Although there are a few examples of stochastic dynamic CGE models (e.g. Manne & Richels 1995a), this work is unique in several key ways. First, it is focused on different questions—while previous work focused on optimal emissions paths/policies over time, this work treats policy as an uncertainty and focuses on optimal electricity investment strategies. Second, this work considers two different uncertainties as well as how they interact. Third, a decision-dependent uncertainty is represented using a novel formulation of stochastic technological learning. Fourth, this work capitalizes on the CGE nature to consider economy-wide impacts of sectoral strategies and policies. Finally, it demonstrates the value of considering uncertainty—illustrating how uncertainty affects investment strategies as well as the expected cost of future policy.

With this problem in mind and the stochastic dynamic CGE framework in hand, this dissertation investigates the following questions, focusing on the United States as a case study:

(1) *Given the uncertainty in future climate policy, what electricity generation and emissions reduction investments should be made in the near-term?*

(2) *How does the optimal near-term investment strategy change when future policies include a safety valve which allows an emissions cap to be exceeded but at a cost?*

(3) *How do uncertain technology costs affect near-term electricity technology investment and emissions reduction decisions?*

For these questions, this work also examines how the optimal near-term decisions under uncertainty differ from the deterministic strategy, thereby demonstrating the importance of realistically representing uncertainty in decision support models.

1.3 Dissertation Structure

This Dissertation is structured as follows. Chapter 2 provides background and context for the current work. It presents an overview of the U.S. electric power sector and its role in producing emissions that contribute to climate change. It also provides background on electricity technology costs and the U.S. policy landscape, as well as the uncertainty involved in each. Chapter 3 provides a review of the literature on decision support models for electricity decisions, methods for studying decision making under uncertainty, and how uncertainty methods have been applied to CGE models.

Chapter 4 describes the numerical modeling framework developed and used to study optimal electricity technology and emission reduction strategies. This modeling framework is then utilized in Chapter 5 to investigate electricity investment strategies under uncertainty in future climate policy. First, cases in which the policy is known for certain are explored, then cases with policy uncertainty. Optimal strategies are explored over a range of expectations about future climate policies. The stochastic strategies and costs are then compared to the deterministic ones. Sensitivity analysis on assumptions about how rapidly low-carbon generation can expand is also included.

Chapter 6 investigates the impact of uncertainty in the costs of technologies on near-term strategy. The future costs of low-carbon technologies are uncertain, but depend on how much we invest in those technologies in the near-term. This decision-dependent uncertainty captures issues of technological learning in a stochastic setting.

The consideration of future policies including a safety valve is explored in Chapter 7. The safety valve presents the option to go over the emissions cap, but at a cost. How that option affects optimal near-term strategy is investigated.

Finally, the key insights, implications and contribution of this dissertation are summarized in Chapter 8. This chapter also includes discussions of limitations and future research opportunities.

Chapter 2: Overview of the U.S. Electric Power System

The focus of this work lies at the intersection of the U.S. electric power system, the climate system and the policy system. At the most fundamental level, almost all production of electricity requires some form of energy to spin a coil in a magnetic field.[1] That energy can come from fossil fuels like coal and natural gas, renewable resources like wind, or other energy sources like nuclear power. Depending on the energy source used, carbon emissions are produced by the generation of electricity. These emissions contribute to global climate change and affect the ability and cost of meeting long-term climate goals. This work focuses on the evolution of the U.S. electric power system in terms of electricity generation technologies and resulting emissions. This chapter provides a background of the relevant structure and activities of the electric power system as well as the landscape of policies that affect the system in the U.S.

2.1 Electric Power as a Complex System

Electric power systems are complex sociotechnical systems. The power generation sector is part of an infrastructure system comprised of the power plants that produce electricity; the homes, commercial buildings and industrial facilities that demand electricity; and the transmission and distribution lines that move supply to demand. Power generation fits within a larger socio-technical system comprised of the physical infrastructure components; the firms and public utilities that own and operate the infrastructure; customers that use the electricity; regulatory agencies that oversee operations and shape market structures, and other stakeholders such as the natural gas and coal industry and environmental non-governmental organizations that have vested interests in this sector. The electricity system is also an important component of a country's wider economy, interacting with other sectors through input and output prices.

In making decisions about how to operate existing electric generation capacity and invest in new capacity, three unique characteristics of electricity must be considered. First, physical laws dictate where electricity flows. We can control how much electricity is produced, how much we consume, and the transmission and distribution network we build, but the precise flows of electricity are governed by Kirchhoff's laws. As a result, siting capacity in relation to existing

---

[1] Solar photovoltaic is an exception in that there is no turbine involved.

infrastructure and load (demand) centers is very important. Second, if demand for electricity is not met the result is power failures, blackouts or system damage. Avoiding these problems requires that generation capacity remain ahead of demand. Third, unlike other commodities, at present electricity cannot be effectively stored on a large scale. While batteries and pumped hydro are storage options, they are limited by the scale at which they can store power and by their high costs. Fuel cells are an emerging technology that could potentially provide large-scale storage on the electricity grid, but they are very high cost and not yet ready for commercialization. The lack of large-scale, cost-effect storage options requires that electricity be produced in "real time" to meet demand. From an operations standpoint, this creates the need to maintain a "spinning reserves" margin above current demand. This also suggests the value of an electricity mix that includes flexible technologies that can start up, ramp up and shut down quickly. These unique characteristics create unique challenges in the electric power sector.

2.2 U.S. Electric Power System

The U.S. electric power system is among the most diverse in the world due to its size, assortment of regulating authorities, and mix of generation capacity. Figure 2.1 shows electricity generation facilities in the U.S., with the size of the circle reflecting relative electricity output. In general, capacity aligns with population. The U.S. power system is currently dominated by dispatchable centralized generation. Dispatchable generation refers to power plants that can adjust their output (e.g. turn on or off or ramp up or down) on demand in a relatively short amount of time (seconds to a few hours). Plants run by renewable energy sources which are variable and intermittent (such as wind and solar) are not dispatchable. Centralized means large facilities that benefit from economies of scale (e.g. coal, natural gas, or nuclear plants, or large-scale wind or solar farms) and transmit electricity over potentially long distances to end users. This is in contrast to distributed generation which refers to electricity produced onsite or close to end users where production capacity is sized towards electricity demand at the site, thereby minimizing transmission. Over time, however, it is possible that the system could include an increasing share of non-dispatchable, distributed generation (i.e. renewables), depending on how policies and technology costs evolve.

This section presents an overview of the U.S. electric power generation sector. The current generation mix, the actors involved, the decisions being made and the role of technology costs in generation decisions are briefly discussed.
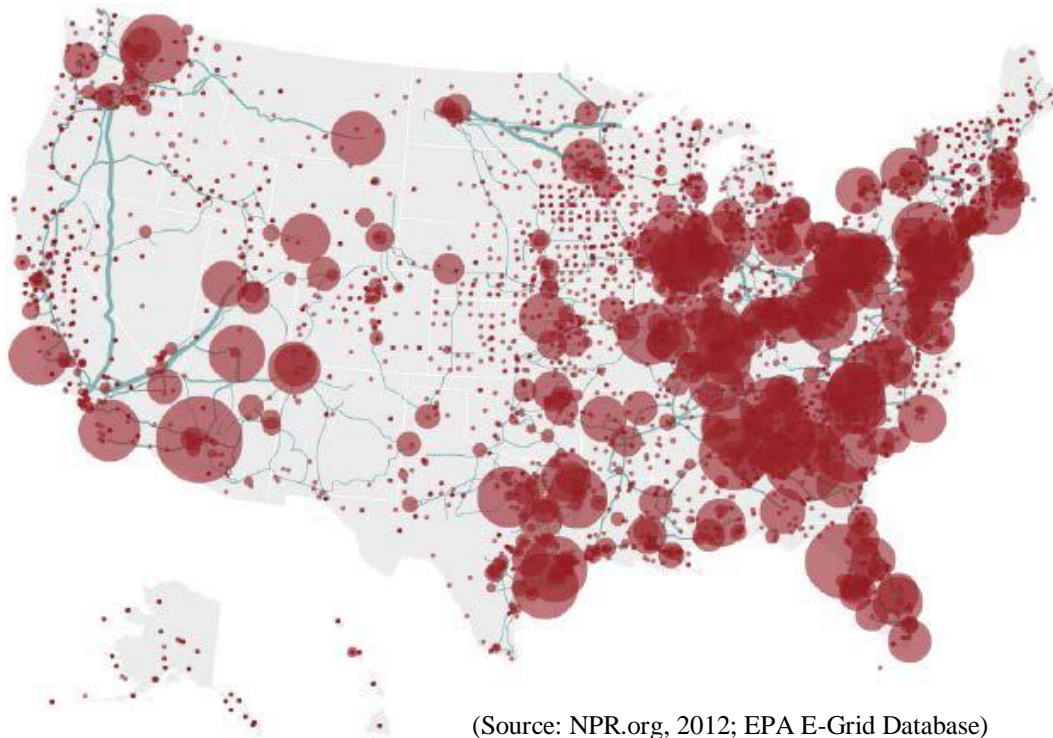
Figure 2.1 Electricity Generators in the U.S.

2.2.1. U.S. Capacity and Generation

The U.S. uses more electricity than every country other than China. In 2011, over 4 billion (4,100,656,000) megawatt-hours (MWh) of electricity were generated by the 18,530 generating units comprising the country's 1,153,149 megawatts (MW) of generating capacity (EIA, EPA2011). That capacity is dominated by natural gas (41% in 2011) and coal (30%). Much of the remainder is nuclear and hydro power (9% each), while renewable sources such as wind, solar and biomass together comprise 5% (Figure 2.2). Of the renewables, wind capacity has surged over the last 5 years. As total electricity demand continues to increase over time (Figure 2.3) and older plants are retired, new capacity must be added to the system.

In terms of actual electricity generation (operation of the installed capacity), coal dominates with 43% of generation in 2011. Natural gas (24%) and nuclear (20%) generation are

the next largest sources. Hydro makes up 8%, renewables 4%, and oil 1% (Figure 2.4). Coal and nuclear make up a higher share of generation than capacity because they are mainly baseload technologies, meaning they operate nearly year-round except for scheduled maintenance. These technologies tend to be baseload because they have high capital costs and relatively low operating costs. Natural gas, on the other hand, makes up a lower share of generation than capacity. The excess natural gas capacity is used to meet peak demand. The existence of excess natural gas capacity means that share of natural gas generation can easily increase under the right conditions. In fact, the share of natural gas generation equaled that of coal generation for the month of April in 2012 when gas prices were exceedingly low (EIA, 2013c). In general, the U.S. generation mix is similar to the global mix—the main difference being that globally hydro makes up 17% as opposed to 8% in the U.S. (and each other source contributes slightly less than in the U.S) (IEA, 2012).

(Source: EIA, 2013b)

Figure 2.2 U.S. Electricity Capacity Portfolio 2011

(Source: EIA, 2011)

Figure 2.3 U.S. Electricity Demand 1949-2011

(Source: EIA, 2011)
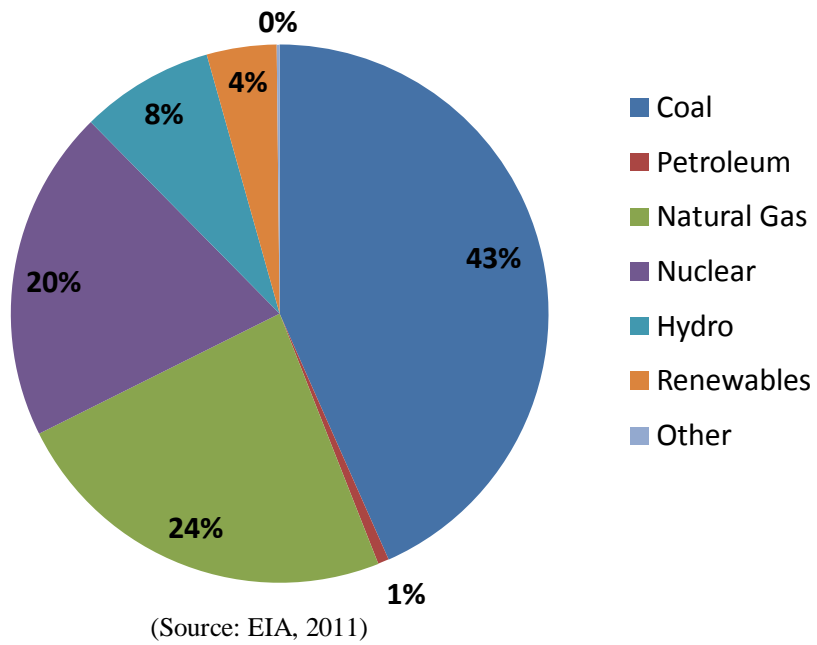
Figure 2.4 U.S. Electricity Generation 2011

2.2.2 Actors in Electric Power

For electric generation, the two main actors involved are the power producers which own and operate generation capacity and the regulatory bodies that try to shape generation. Of course, generation is also affected by consumers (residential, industrial and commercial) demanding power and by other actors, such as the coal and natural gas industries which affect resource availability and fuel costs.

Power producers come in several forms. Investor-owned private utilities produce over 65% of the electricity consumed in the U.S. on a kilowatt hour basis. Comprising only about 6% of utility entities, these are typically large, vertically integrated (owning and operating generation, transmission, and distribution facilities) and regulated. Other utilities are publicly owned (such as municipal utilities), federally owned (such as large hydroelectric projects), or are rural cooperative utilities. As a whole, producers are making decisions about how to expand electric generation capacity.

Utilities are connected to a large and diverse power system that is organized and operated differently by region. The U.S. is divided into three main power grids: the Eastern Interconnect, the Western Interconnect and the Texas Interconnect (Figure 2.5), with limited interconnections to each other. The Western and Texas Interconnects are linked to Mexico and the Eastern and Western Interconnects are strongly connected with Canada. Figure 2.5 also shows the ten North American Electric Reliability Councils (NERC regions), which are structured to oversee reliability and security of generation resources and electricity supply. Utilities within each NERC region coordinate planning and operations in order to maintain security and reliability of supply. The system is also divided into a number of Regional Transmission Organizations (RTOs) (or Independent System Operators (ISOs)) (Figure 2.6). RTOs exchange power and coordinate planning and operation between regional grid areas.

Oversight of electric generation is the responsibility of several regulating bodies at different levels of government. At the federal level, the Federal Energy Regulatory Commission (FERC) is the main regulating body, responsible for the oversight and approval of wholesale electricity and transmission rates in interstate commerce. Also at the federal level, the Environmental Protection Agency (EPA) administers most environmental regulations that affect generation facilities, such as emission and technology standards for power plants. The Nuclear Regulatory Commission (NRC) is another federal agency that is in charge of nuclear power

safety regulation. At the state-level, the main authorities are Public Utility Commissions (PUCs) which are responsible for regulating the rates and services of publicly-owned utilities. State environmental agencies are also involved by administering federal regulations or enacting state-level regulations.

Overall, the U.S. electric power system is a hybrid of regulation and competition. In the early 1990s the U.S. began moving toward a deregulated system that would allow competition among electricity producers while generally retaining regulation of retail rates.[2] A number of states pursued restructuring toward deregulated markets; however some have since suspended their efforts. The result is a mix between highly regulated monopolies and liberalized wholesale generation markets. These market differences can affect generation capacity expansion decisions.
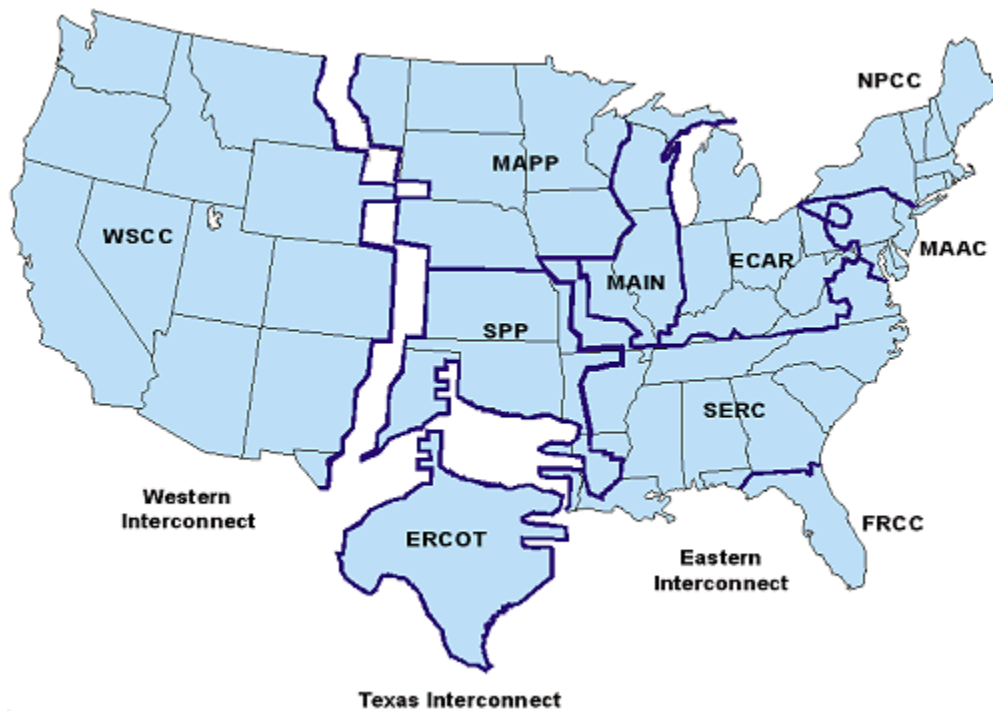
Figure 2.5 U.S. Power grids and NERC Reliability Regions

---

[2] In the electric power sector, deregulation applies just to wholesale generation, not other components of the system like transmission and distribution
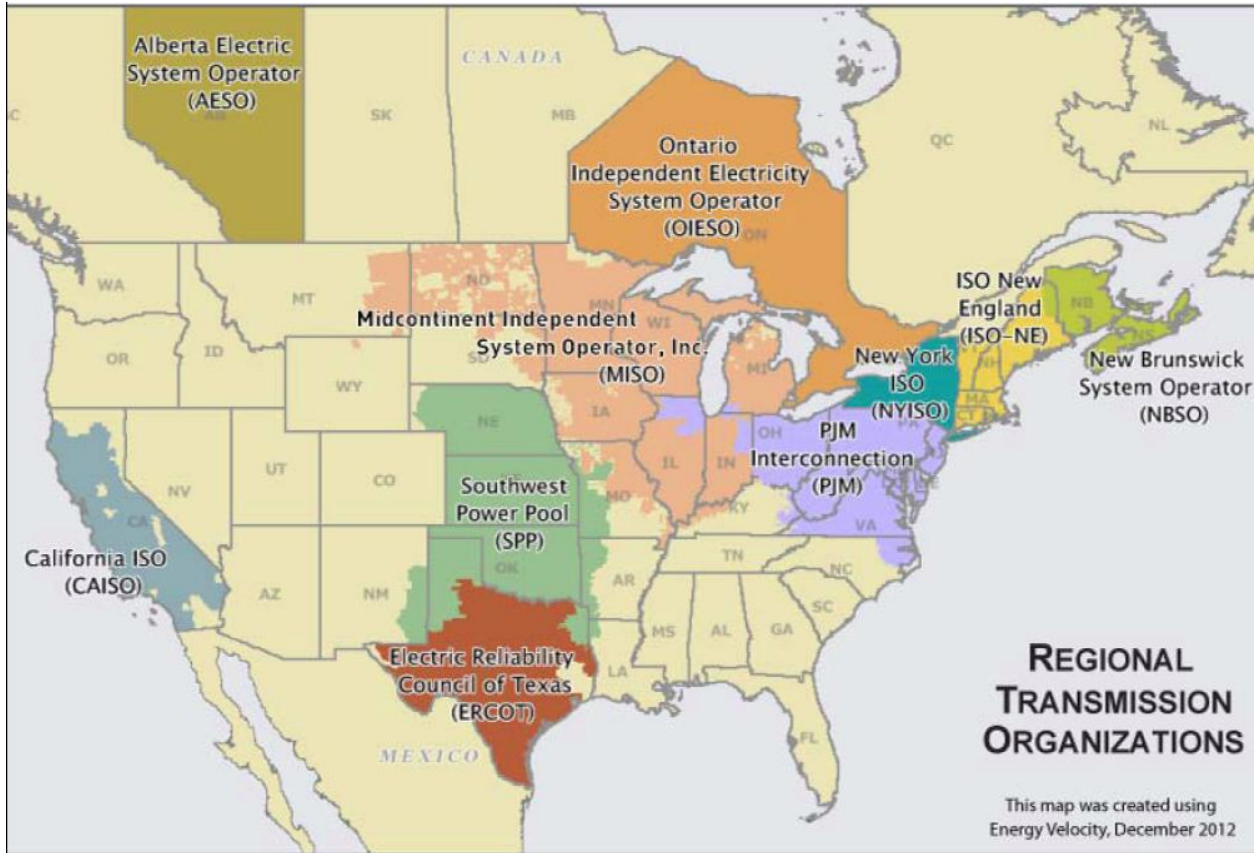
Figure 2.6 Regional Transmission Organizations in North America

2.2.3 Electric Power Decisions

Within the complex and diverse electric power system there are a number of
interdependent decisions involved at different levels and timescales ranging from real-time (or
seconds) to 30-50 years into the future (Figure 2.7). Economic dispatch, the decision about how
to operate existing generation capacity to best meet demand, takes place on the hourly time scale.
Generator control management takes place at the seconds to minutes scale to meet fluctuations in
demand and keep the system operating reliably. Unit commitment, which proceeds on the order
of hours up to a week in advance of actual generation, consists of operators of the system and
electric utilities confirming that specific generating facilities will be available to generate
electricity during the next period in question. Electricity demand forecasting is needed by system
operators to be able to plan for the required growth in generating capacity. At the longest time
scale (years to decades) is the task of generation expansion planning, which entails making
decisions about the amount and type of generation capacity to add to the system. There are two

additional long-term decisions affecting the system that are not shown in the figure: transmission planning and regulation. Transmission planning decides how to expand transmission corridors in order to handle increased generation capacity and deliver power reliably. Finally, regulation, which involves decisions about how power markets are configured and rules that are enforced (such as emission limits on power generation), directly impacts operation and expanssion decisions.

(Source: Perez-Arriaga et al., 2009)
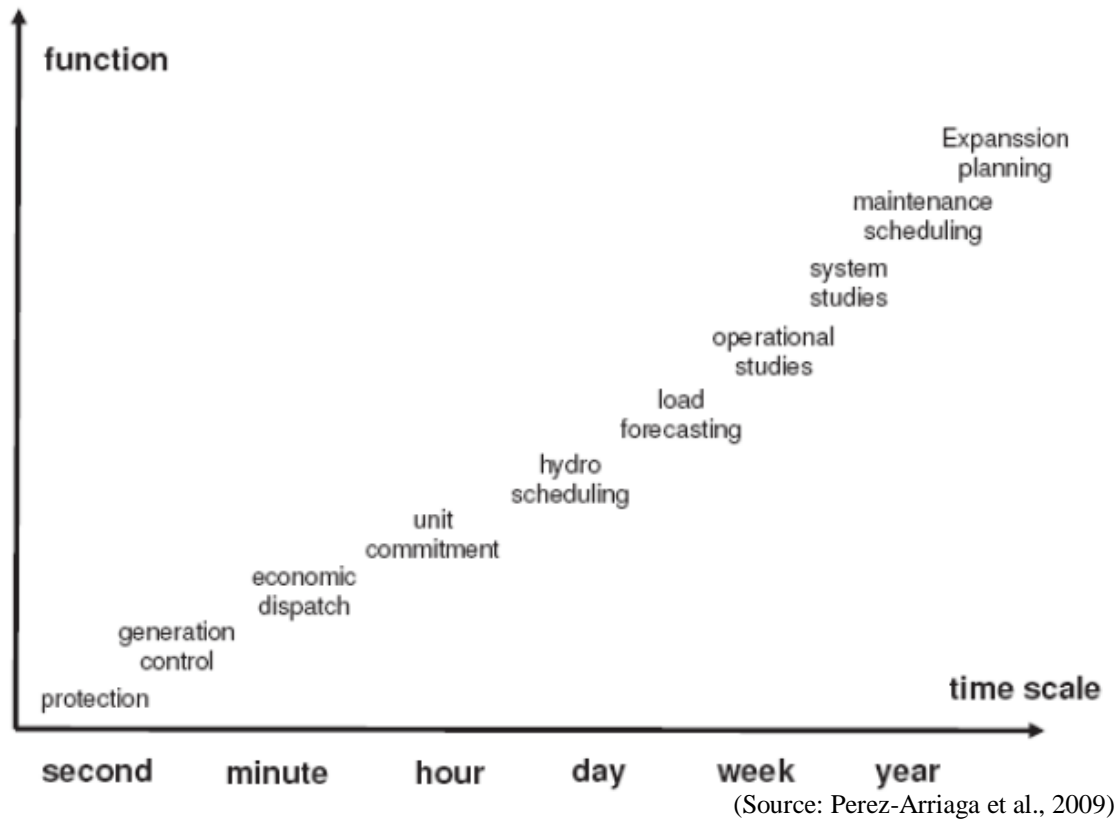
Figure 2.7 Electric Power System Planning Functions

2.2.4 Technology Costs

Decisions about which generation technologies to build and operate are ultimately driven by costs. Levelized cost of electricity (LCOE) is one common metric used to compare the costs of the alternatives. LCOE is the price of electricity per kilowatt-hour necessary to recover all of the expenses over the life of a plant, such that the net present value of the plant is zero. Table 2.1 shows the latest LCOE estimates from Morris *et al.* (2010). According to this study coal and

25

natural gas generation are very close in cost and wind is the most competitive advanced (zero or low-emitting) technology. However, these numbers are only one example, and the range of LCOE estimates for generation technologies is very wide.

Table 2.1 Levelized Cost of Electricity Estimates

	Levelized Cost of Electricity ($/kWh)
Coal (Advanced PC)	0.054
Coal with CCS	0.092
Natural Gas Combined Cycle	0.056
Natural Gas Combined Cycle with CCS	0.085
Nuclear	0.088
Biomass	0.085
Wind (Onshore)	0.077
Solar Thermal	0.194
Solar Photovoltaic	0.290

Also crucial to the overall cost of technologies and the decision of what to build is the relative contribution of fixed and variable costs. Table 2.2 shows EIA estimates of capital costs, fixed operations and maintenance (O&M) costs, and variable O&M (EIA, 2013a), as well as a 5-year average of fuel costs from 2005-2010 calculated from EIA data. Variable costs (variable O&M and fuel costs) increase with output while fixed costs (capital and fixed O&M) do not. Although these numbers are also just one example of a range of estimates, they can provide intuition to the roles of various cost components. Coal, coal with CCS, nuclear, solar, and biomass are the most capital-intensive plants, while also having relatively low variable O&M and fuel costs. Natural gas plants have low capital costs, but high variable O&M and fuel costs. However, recent success in extracting shale gas in the U.S. has lowered natural gas prices to be competitive with coal. Wind and hydro are on the lower side of capital costs and have zero variable costs. When combined, these fixed capital and variable costs are used to determine the least-cost way to meet electricity demand and explain why a diverse portfolio of technologies is often most economical.

Table 2.2 Capital, Fixed O&M, and Variable O&M Cost Estimates

	Overnight Capital Cost ($/kW)	Fixed O&M Cost ($/kW-yr)	Variable O&M Cost ($/MWh)	Variable Fuel Cost ($/MWh)[1]
Coal (Advanced PC)	$3,246	$37.80	$4.47	$10.43
Coal with CCS	$5,227	$80.53	$9.51	$11.63
Natural Gas Combined Cycle	$1,023	$15.37	$3.27	$38.50
Natural Gas Combined Cycle with CCS	$2,095	$31.79	$6.78	$45.56
Natural Gas Conventional Combustion Turbine	$973	$7.34	$15.45	$63.54
Nuclear	$5,530	$93.28	$2.14	$6.62
Biomass	$4,114	$105.63	$5.26	$8.02
Wind (Onshore)	$2,213	$39.55	$0.00	$0.00
Solar Thermal	$5,067	$67.26	$0.00	$0.00
Solar Photovoltaic	$4,183	$27.75	$0.00	$0.00
Geothermal	$4,362	$100.00	$0.00	$0.00
Hydroelectric (Conventional)	$2,936	$14.13	$0.00	$0.00
Hydroelectric (pumped storage)	$5,288	$18.00	$0.00	$0.00

Source: EIA (2013b)
Note: Costs in 2012$
[1]Because fuel prices can be volatile, here they are calculated as a 5-year average from 2005-2010 using EIA data on fuel prices and heat rates for generation technologies (efficiency of turning fuel energy into electricity)

Plants with high capital costs but low variable costs (such as coal, nuclear and hydro) are ideal for meeting baseload demand (the base level of demand that is guaranteed at all times). Once the expensive investment is made to build such a plant, it is cheap to operate and should be run nearly all the time. Additional demand should be met by technologies that have lower capital costs even if relatively high variable cost, since they will only be operated some of the time, when demand and prices are high. Natural gas turbines which are cheap to build but expensive to operate (but quick to turn on) are typically used to meet "peak" demand (short periods of time when demand, and therefore costs, are highest). Intermediate "shoulder" demand has largely been met by NGCC plants, as well as older coal and nuclear plants. However, NGCC is increasingly being used for baseload, driven by low fuel costs due to shale gas. Renewables are unique because the availability of the resource, particularly for wind and solar, is an additional concern. Although wind generation is fairly competitive cost-wise, the intermittency of the wind

resource limits wind's ability to be adopted at large-scale, in the absence of economically storage. Renewable generation that is built is usually utilized whenever available, and the rest of the system adjusts.

2.3 U.S. Energy Policy Landscape

While the previous section discussed actors involved in the regulation of the electric power sector, this section focuses on broader government energy policies that impact electric power. Such policies are important in shaping electricity generation decisions. This section provides an overview of the electric power sector's role in climate change and the roles and types of government intervention.

2.3.1 Climate Change and Electric Power

Climate change poses an ever-increasing risk as the emissions of carbon dioxide (CO_2) and other greenhouse gases (GHGs) rise. Emissions, mainly from the combustion of energy sources, have been steadily increasing since pre-industrial times (Figure 2.8). There has been a slight decrease in emissions over the last several years due to decreased economic activity as a result of the global economic recession. However, from a climate perspective, the total stock or concentration of emissions in the atmosphere is what matters. The yearly flows of emissions released into the atmosphere far outweigh the yearly sinks that remove emissions from the atmosphere. As a result, even though yearly emissions have decreased in the last few years, total concentrations of emissions continue to rise (Figure 2.9). Presently concentrations of CO_2 emissions in the atmosphere are approximately 400 parts per million (ppm). The U.S. was the largest emitter in the world until 2006 when China surpassed the U.S. However, the U.S. remains the largest cumulative contributor of CO_2 in the atmosphere.

(Source: EIA, 2011)

Figure 2.8 Total U.S. CO_2 Emissions

(Source: NOAA, 2013)

Figure 2.9 U.S. CO_2 Concentrations

This work focuses on the electricity sector and its emissions. The electricity system directly contributes to climate change through the emission of greenhouse gases during the combustion of fossil energy (coal, oil, and natural gas) to generate electricity. The electric power sector was responsible for approximately 40% of U.S. CO_2 emissions in 2012, or approximately 2000 million metric tons (EIA, 2013a). The transportation sector is the next largest emitter at 34%, with other sectors (commercial, residential, industrial) making up the remaining 27% (Figure 2.10).

Different energy sources have different carbon contents and therefore produce different amounts of emissions. Coal is the most carbon intensive, followed by oil and then natural gas. The high carbon-intensity of coal and its high share in electricity generation make coal responsible for nearly 80% of electric power sector emissions (Figure 2.11). Natural gas accounts for the overwhelming majority of the remaining emissions from the electric power sector.

(Source: EIA, 2013b)

Figure 2.10 CO_2 Emissions by Sector 2011

Figure 2.11 CO_2 Emissions from the Electric Power Sector by Fuel Source 2011

To reduce electricity emissions, society must shift to cleaner energy sources for generation and/or reduce overall energy use. The second option could mean reducing total electricity consumption or increasing the efficiency of generation to use less energy for the same amount of output. It is important to keep in mind that because of the long lives of power plants, new capacity investment decisions today are linked to operation decisions for decades to come, and emissions are the net result of both of these decisions.

While the electricity system affects the climate system through emissions, at the same time the climate system can impact the electricity system. Climate directly impacts renewable energy sources such as hydropower, wind, solar, and biomass (e.g. Wang and Prinn, 2010). Taking hydropower as an example, a warming climate can lead to less precipitation and more evaporation leaving fewer water resources for the generation of electricity. In addition, fear of climatic change can also steer society away from the most carbon intensive energy forms like coal and oil. We might expect to see that as climate changes with rising global average temperatures, people will become more willing to change their energy usage and switch to cleaner forms. A key dimension in the relationship between the electricity system and the climate is the time delay. It takes time for emissions to change climate and time to see the impacts of a changing climate. The challenge is to plan the electricity system to hedge against the risks of climate change before we see or fully understand those risks. While climatic impacts and risks

31

are of great importance, they are beyond the scope of this work, which focuses on the economic costs and risks associated with climate policy and electricity technology investments.

2.3.2 Role of Government

Government intervention in the electricity sector is often justified as a way to correct market failures, reduce market barriers and stimulate new clean energy technologies. Three main arguments are put forward.

First, is the argument that the government needs to correct the market failure of negative environmental externalities caused by the combustion of fossil fuels. The burning of coal, oil, and natural gas for electric power currently accounts for about 40% of CO2 emissions, more than 26% of smog-producing nitrogen oxide emissions, 33% of toxic mercury emissions, and 64% of acid rain-causing SO_2 emissions (Nogee *et al.*, 2007). These emissions have harmful effects on the climate, environment, and human health. These negative externalities constitute a market failure that suggests government policy intervention. An array of economic work supports broad incentive-based measures, such as a cap-and-trade system or emissions tax, over technology specific measures for addressing environmental externalities such as climate change (for example, Baumol and Oates, 1988; Tietenberg, 1990; Stavins, 1997; Palmer and Burtraw, 2005; Dobesova *et al.*, 2005). By directly targeting the source of the problem (the emissions), policies that place a price on emissions internalize the externalities

The second argument for government intervention is based on the public good nature of knowledge which creates a market failure in which there is an underinvestment in research, development and deployment. The benefits of RD&D cannot be fully captured by the private entity that undertakes the investment. The knowledge gained through such investment becomes available to everyone, and the imperfect patent system prevents the investor from capturing all of the gains. Even if the patent system worked perfectly, there would still be knowledge spillover effects in which knowledge gained from investment and learning in one thing can be used to benefit another thing. As a result, competing firms are likely to benefit from the innovation without contributing to the development cost (Mowery & Ziedonis, 2001). These positive externalities are not factored into investment decisions and as a result there is an underinvestment in RD&D by the private sector according to what is optimal from a social welfare perspective. It can then be argued that government technology investments are needed to

solve this problem and incentivize people to invest in clean technologies, or alternatively do the investing itself. Otto and Reilly (2008) investigated the need for policies targeting specific low-carbon technologies. They found that when such technology externalities exist, adoption or R&D subsidies added to a CO_2 trading scheme can increase the cost-effectiveness of achieving an abatement target by internalizing the externalities. They found directed R&D subsidies do better in terms of welfare than adoption subsidies. However, they noted that depending on the target, a CO_2 trading scheme alone can be sufficient to induce adoption of low-carbon technologies, alleviating the need for technology specific policies.

The third, and related, argument for government intervention focuses on learning and scale effects, which act as market barriers for new technologies. New technologies, like renewables or CCS, have to compete with established technologies which have benefited for a long time from scale, mass production and learning effects, all of which lower costs. When new technologies arrive on the market, they have not reached an ideal level of performance in terms of cost and reliability. Optimum performance is achieved gradually through learning by doing or learning by using (Arrow, 1962; Dosi, 1988). There are two forces at work: a technology may be adopted because it is efficient or it may become efficient because it is adopted (Arthur, 1989). Because new technologies are not economically competitive with established fossil technologies, they may not be adopted and remain underdeveloped, producing a "lock-in" phenomenon. Although commercialization is typically viewed as the responsibility of the private sector, it can be argued that in order to save new clean technologies from the "valley of death" between development and deployment, government involvement is needed in the emergence phase to protect new technologies from direct competition with conventional technologies (PCAST, 1997). Doing so will increase their adoption, allow them to progress along their learning curves, and create economies of scale, thereby reducing the costs of the technologies. Without such support, market forces alone would likely result in only limited diffusion of new technologies in a few market niches. However, "picking winners" can be a difficult and costly strategy. It is difficult to determine when early investment is warranted to help a potentially valuable technology and when it is wasted resources on a bad technology. Some technologies just may not be well suited competitors for the market and it would be unwise to spend resources trying to protect them.

2.3.3 Types of Government Interventions

Given that government intervention in the electricity sector can often be justified by the market failures and barriers described above, there are different forms that an intervention can take. Government sponsored research and development (R&D) can create new technologies, technology improvements and reductions in cost. Demonstration projects can contribute to technology development by building a customer base and providing some experience with the technology, market, and regulations. Tax incentives and subsidies help reduce technology costs and encourage greater market penetration or further investment and development. Environmental policies and regulations can correct negative externalities while also creating or stimulating a market for advanced low-carbon technologies.

Research and Development

Government funded R&D, the classic "technology-push" mechanism, is intended to overcome the private sector's underinvestment in R&D. The federal government offers direct support through R&D contracts with private firms, contracts with grants and universities, research conducted within government agencies, and contracts with industry-led consortia or collaboratives. R&D as a whole is intended to develop technology and lower costs.

The U.S. has a large budget for energy-related R&D. Substantial support came from the American Recovery and Reinvestment Act (ARRA) of 2009, where $90 billion was provided for clean energy investments. Out of that amount $60 billion was specified for projects directed by U.S. agencies and $30 billion was provided by tax incentives. In total, $26.6 billion of ARRA funding was devoted to renewables. Other areas receiving funding include energy efficiency, grid modernization, carbon capture and storage, green innovations, advanced transportation and nuclear. Including ARRA funds, the fiscal year 2011 budget request included over $17 billion for energy technology innovation.

ARRA also created new R&D programs. The flagship is the Advanced Research Projects Agency-Energy (ARPA-E) program, which was initially funded with $400 million. ARPA-E supports research that has high potential commercial impacts but is deemed too risky for industrial investments. ARPA-E does not fund discovery science or incremental improvements to current technologies. The funding range per project can range from $500,000 to $10 million. Projects are selected on their potential to make rapid progress towards commercialization and

will not be extended without demonstrable progress in a 2-3 year timeframe. ARRA also created Energy Innovation Hubs, each of which comprise a large set of investigators spanning science, engineering, and policy disciplines focused on a single critical national need identified by the DOE. So far there are three hubs: Fuels from Sunlight, Efficient Energy Building Systems Design, and Modeling and Simulation for Nuclear Reactors. Additionally, newly formed Energy Frontier Research Centers advance fundamental science relevant to real-world energy systems. Each focuses on the long term basic research needed to overcome roadblocks to innovative energy technologies in a particular area. This research is fundamental science motivated by the need to solve a specific problem, such as energy storage, photoconversion, CO_2 sequestration, etc.

The DOE also expends substantial resources on its nuclear energy R&D program which includes advanced reactors, fuel cycle technology and facilities, and infrastructure support. Nuclear research programs include Generation IV Nuclear Energy Systems which focuses on advanced reactors not yet close to deployment and the Fuel Cycle Research and Development Program which focuses on engineering-scale and prototype reprocessing facilities.

Demonstration

Failure to develop a technology by demonstrating its role in the future energy system "send[s] the wrong signal to governments and private investors and will have the potential to inhibit future investment" (RETD, 2006). As such, governments are increasingly turning to demonstration projects. The demonstration phase is an excellent learning experience for industry because it allows feedback from customers, engineers, and installers that can be integrated into product design. Small changes and details to tailor the technology to the market and regulations reinforce the product development and the technology may also be improved or become cheaper through learning. For investors, government-supported demonstration can be an excellent way to reduce risk. The chance to observe the technology performing will increase the information they have about the technology's shortcomings and successes.

In the U.S. electric sector government demonstration has focused on CCS. $1 billion of the ARRA stimulus funding was awarded to the FutureGen Alliance to build FutureGen 2.0, a clean coal repowering program and carbon dioxide (CO_2) storage network. Also, the DOE has

set up a grant of $575 million to fund 22 industrial CCS projects in 15 states. However, most of these projects have stalled.

Tax Incentives and Subsidies

Another common form of government intervention is through tax incentives and subsidies. These programs help reduce costs and encourage greater market penetration or further investment and development.

The main tax incentives for energy projects are the production tax credit (PTC) and the investment tax credit (ITC). The Renewable Electricity Production Tax Credit is based on actual energy production by eligible technologies. Initially enacted in 1992, it currently provides 2.2 cents per kWh for generation from qualified wind, geothermal and closed-loop biomass facilities, and 1.1 cent per kWh for other eligible technologies. Between 1992 and 2007 the PTC experienced several lapses in funding, creating policy uncertainty and significantly affecting the actual adoption of technology (Figure 2.12).
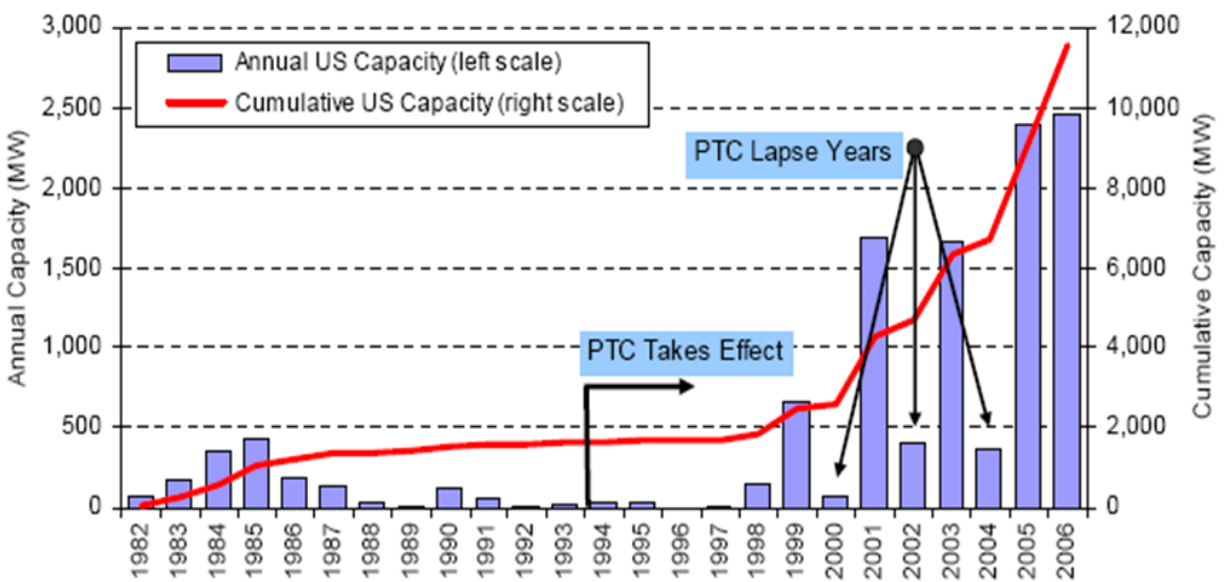
(Source: Wiser et al., 2007)

Figure 2.12 Impact of Production Tax Credit on Wind Power Capacity

The federal Business Energy Investment Tax Credit is available for solar facilities, geothermal energy property, geothermal heat pumps, fuel cells, combined heat and power (CHP, or cogeneration facilities), microturbine plants, and small commercial wind energy property (100 kW or less). The ITC is a dollar-for-dollar credit against tax liability, claimed entirely in the year

36

the project is placed in service. The ITC is 30% of the capital cost for solar, fuel cells and small wind, and 10% for geothermal, microturbines, and CHP. An ITC is also available for power generation projects that use IGCC or other advanced coal-based electricity generation technologies that include equipment that separates and sequesters 65 percent of the project's total CO_2 emissions.

Additional tax incentives include those for research expenditures and the Modified Accelerated Cost-Recovery System (MACRS). The research and experimentation (R&E) tax credit is 20% of qualified research expenses above a base amount determined by the taxpayer's historical research intensity. MACRS applies to all ITC-eligible projects and allows them to be depreciated on an accelerated (200% or double declining balance) basis over five years.

Another form of subsidy are DOE loan guarantees, which are made available under Title XVII of the Energy Policy Act of 2005 (EPAct05) to projects that avoid, reduce, or sequester air pollutants or anthropogenic emissions of greenhouse gases. The final regulation provides that the DOE may issue guarantees for up to 100 % of the amount of a loan, subject to the EPAct05 limitation that DOE may not guarantee a debt instrument for more than 80 % of the total cost of an eligible project. A broad range of projects are eligible for funding, including renewables, CCS, and nuclear, among others.

There are several federal programs that aim to reduce the costs of nuclear generation. EPAct05 allows standby support, or regulatory risk insurance, to help pay the costs associated with regulatory delay at up to six commercial reactors. The Nuclear Power 2010 program encouraged near-term construction of new commercial reactors by providing up to half the costs of licensing plant sites and reactors and preparing detailed reactor designs. The DOE also provides licensing and engineering assistance to small modular reactor designs. Nuclear accident liability is given by the Price-Anderson Act, and extended by EPAct05.

Environmental Policies and Regulations

There is a wide variety of environmental policies and regulations in the U.S. that intended to correct the negative externalities associated with emissions. Most are command-and-control regulations that directly limit emissions, often through technology performance standards. For example, the Clean Air Act sets National Ambient Air Quality Standards for six criteria pollutants (carbon monoxide, lead, nitrogen oxide, ozone, particulate matter, and sulfur dioxide). Other command-and-control regulations set technology standards that must be met, for example

requiring scrubbers in the stack pipes of coal plants. A series of court decisions and budget appropriations have given the Environmental Protection Agency (EPA) greater authority and funding to impose much greater levels of regulatory restrictions on power plants than in the past. EPA's new and proposed emissions regulations have focused on coal power plants.

Another technology-specific policy is a renewable portfolio standard (RPS), which requires that a certain percentage of electricity be produced by renewables. Currently 29 states and the District of Columbia have an RPS mandate and five additional states have voluntary renewable goals. There have also been several proposals for a national RPS or Clean Energy Standard (which expands the RPS concept to other technologies like nuclear, CCS and efficiency), such as the American Clean Energy Leadership Act (S. 1462) and the Practical Energy and Climate Plan Act (S.3464). The government also uses procurement to require the use of a certain technology throughout federal government operations.

As discussed previously, economists consider an emission pricing scheme, such as an emissions tax or cap-and-trade system, the most cost-effective way to reduce emissions. An emissions tax places a tax on each unit of emissions produced. A cap-and-trade system places a limit on the amount of emissions, issues permits for each unit of emissions allowed, and allows trading of those permits which creates a market and therefore a price on the permits and the unit of emissions they represent. Fischer and Newell (2004) compared the partial equilibrium social cost of different policies using a simple economic model of electricity markets. They found that a renewable portfolio standard (RPS) set to achieve a given reduction in CO_2 emissions is 7.5 times as costly in terms of social welfare as using an emissions tax (assumed equivalent to a cap-and-trade) to achieve the same emissions reductions. By shifting investment away from the least-cost emission reduction options and toward specific renewable technologies, which are not necessarily least-cost (or even low-cost), an RPS adds to the economy-wide cost of the policy. A carbon pricing policy does not attempt to pick winning technologies—by forcing fossil fuels to internalize the cost of their emissions, a carbon pricing system indiscriminately provides an advantage to technologies in proportion to the level of emissions they produce, and lets the market choose the least-cost options that achieve the emissions goal.

The notable example in the U.S. of a federal market-based environmental policy is the SO_2 cap-and-trade system under EPA's Acid Rain Program. In recent years, there has been a surge of legislative proposals at the federal level that would put a price on GHG emissions. In

June 2009, the House of Representatives passed the Waxman-Markey bill (the American Clean Energy and Security Act, H.R.2454), which would have established a cap-and-trade program for CO_2 and other GHG emissions. Negotiations on corresponding Senate legislation stalled. Kerry-Lieberman (American Power Act), Kerry-Boxer (S.1733), and Cantwell-Collins (CLEAR Act) were the leading contenders. All of these implemented a cap-and-trade policy which would put a price on emissions and therefore incentivize clean alternative technologies. None of these proposals have been implemented. However, there are regional GHG cap-and-trade programs such as the Western Climate Initiative and the Regional Greenhouse Gas Initiative (RGGI, in the Northeast). There is also a possibility that a national GHG cap-and-trade system may be designed and implemented by executive order via the EPA (like the SO2 trading system).

Although the government policies discussed above are very different in their approach, they all help create or stimulate a market for low-carbon technologies. For this reason they are often referred to as "demand-pull" or market stimulation measures in terms of technologies. Command-and-control regulations may necessitate the development and use of new technologies in order to meet the standard. Technology standards, like an RPS, force the adoption of certain technologies. Government procurement creates an instant market that can influence the entire commercialization chain (Virdis *et al.*, 2003). Carbon pricing policies increase the cost of fossil technologies thereby increasing demand for clean alternatives as firms seek to minimize compliance costs.

2.4 Key Challenges of Electric Power Investment Decisions

This work focuses on near-term electricity generation investment decisions, asking what types of new generation should be built in the next 10-15 years. This question is posed from the perspective of central planner. Such a perspective abstracts away from the complex details of the electric power sector and its regulatory structures[3], but can offer valuable insight in electricity generation planning. From the central planner perspective, the goal of generation expansion planning is to add new capacity at the lowest cost while still meeting demand and system objectives such as reliability or sustainability.

---

[3] For a detailed review of the regulation of the electric power sector, see Pérez-Arriaga (2013).

The answer to the question of what type of new generation to build in the near-term depends on several factors, including existing capacity available to be dispatched, demand, transmission availability, technology costs, policies, and the uncertainty surrounding each. Another crucial component of generation expansion is the irreversibility of decisions. Power plants are expensive investments, and once built they are typically in operation for 40-60 years. The capital intensity of power production means investment is based on the assumption that the plant will operate for a number of decades. To the extent possible given uncertainty, we want to avoid making investments in plants that may be underutilized or shut down for economic reasons in the future. As a result, generation expansion, also involves operation of built capacity as a subproblem. In terms of meeting long-term environmental goals, emissions are the net effect of both of these decisions. These complex policy-investment-operation interactions are not yet well understood in policy realms or well represented in decision support tools. This work seeks to provide insight to how to balance expansion and operation decisions in the face of potential emission reduction policies.

To do so, this work focuses on two key challenges affecting electricity investment decisions: (1) uncertainty in future climate policy and (2) uncertainty in the future costs of technologies. The uncertain prospect of future climate policy impacts the solvency of electricity generation investments made today. If a climate policy is implemented during the early or mid lifetime of a power plant, it greatly affects how cost-effective it is to run that plant, which in turn affects how cost-effective it is to build that plant to begin with. Technology costs drive investment decisions and in turn emissions and the cost of reducing emissions, so uncertainty in these costs impact near-term investment decisions. Additionally, investments in cleaner generation today could help reduce the costs of those technologies and therefore the cost of meeting future policy.

Of course there are uncertainties other than technology costs and policy that affect electricity capacity expansion decisions. Some of these include demand growth, fuel prices, infrastructure adequacy, and renewable resource availability. Demand growth is important in determining how much capacity to add to a system. In recent years, uncertainty in demand has increased due to increased electricity conservation and efficiency measures available to consumers. Fuel pricing, driven mainly by resource availability and demand, has also long been a source of uncertainty. While coal prices have historically been low and stable, natural gas prices have been

quite volatile. Recent success in recovering shale gas in the U.S. has resulted in low gas prices. However, long-term price uncertainty will persist. Another uncertainty is whether infrastructure like transmission and distribution line will be adequate for certain expansion plans. A newer uncertainty faced in capacity planning is caused by the variability of renewable energy resources, particularly wind. There are also several uncertainties related to the modeling of the electricity system. For example, elasticities of substitution, labor growth, and energy efficiency improvements are all important model inputs that affect model capacity decisions. While these uncertainties are important and worthy of their own study, this work chooses to focus on uncertainties of technology cost and policy. Costs are ultimately the main driver of technology choice, making cost uncertainty extremely important. Policy directly impacts the cost of technologies and the value of investments, and policy uncertainty represents an important real-world decision-making context.

Given the uncertainty involved, it is important to design an electricity generation expansion plan that hedges against the economic risks associated with policy and technology cost uncertainty in order to provide the best chance of pursuing a least-cost expansion plan. It is the goal of this work to provide insight into such a plan. The complexity of this task suggests the need for decision support tools. The next chapter reviews the literature on state-of-the-art decision support tools for the electric power sector as well as methods for representing uncertainty in such models.

(This page is intentionally left blank)

Chapter 3: Literature Review

Electric power generation decisions are very complex—they must be made within the context of a unique and diverse sector, under uncertainty in future policy and costs of technologies, and considering that investments are quite irreversible due to their high cost and long lifetimes. This complexity presents the need for decision support tools to aid both public and private investors. This chapter reviews the literature on state-of-the-art decision support tools for the electric power sector and methods for representing uncertainty in such models. As this dissertation work employs a computable general equilibrium (CGE) model, this chapter highlights the CGE approach and previous applications of uncertainty.

3.1 Types of Decision Support Tools for the Electric Power Sector

There are different types of models that can help inform energy and climate decisions. Modeling always entails some simplification with the challenge being to simplify those aspect of the world that are not essential to the question being examined, while retaining detail on those things that matter most to the question. This work seeks to investigate near-term electricity investment decisions that maximize expected economy-wide consumption given uncertainty in future policy and technology costs. To do so requires a decision support tool that includes several essential elements: (1) an electric power sector that includes at least conventional and low-carbon generation options and represents the long lifetimes of these capital investments (via capital vintaging); (2) the ability to model climate policies, particularly limits on emissions; (3) an economy-wide framework with multiple sectors that can measure social welfare implications and policy costs; and (4) the ability to represent sequential decision-making under uncertainty with learning.

A main tradeoff in choosing the model type for any analysis is between technological detail and economic completeness. In this regard, there are two general model types: 1) bottom-up engineering cost models and 2) top-down economic models. Bottom-up models are rich in technological detail, but lack economic details and feedbacks from other sectors. On the other hand, top-down models represent microeconomic principles, but often lack technological detail, such as physical engineering constraints.

3.1.1 Bottom-Up Engineering-Cost Models

Bottom-up engineering-cost models use engineering data and principles to represent detailed technical characteristics. With such a high level of technical detail, keeping these models tractable requires a direct cost accounting framework or a partial-equilibrium perspective. A partial-equilibrium model represents one or several sectors in great detail, but does not capture interactions between these sectors and the rest of the economy. As a result, supply, demand, and prices are all exogenous inputs to these types of models. The structure and solution approach of these models varies considerably. Most use a linear-programming or mixed-integer programming optimization framework, or a simulation framework such as system dynamics, agent-based models or game theoretic models (Azar & Dowlatabadi, 1999). Some of the most well-known optimization models for electricity and environmental policy analysis are the MARKAL model (Loulou *et al.*, 2004), the National Renewable Energy Laboratory's ReEDS (Regional Energy Deployment System) model (Short *et al.*, 2009), the International Institute for Applied Systems Analysis (IIASA) MESSAGE model (Messner, 1997), and the EPA's Integrated Planning Model (IPM) (EPA, 2010). These models are specifically designed to study the energy and/or electricity sector. They capture multiple regions, time periods, and technologies (typically 20 or more electricity generation technology types). Of the simulation approaches, agent-based modeling has been useful for short-term electricity decision making, such as operation and bidding decisions, rather than long-term capacity expansion decisions (Sanchez, 2008). The same is true for game theoretic models, though there some good examples of game theory applied to electricity generation expansion (Haikel, 2009; Murphy and Smeers 2005). System dynamics has been criticized in the power system planning community for not being as controlled and transparent a method relative to others (Sanchez *et al.*, 2012).

Bottom-up models typically seek to identify the least-cost method of operating and/or expanding electricity generation technologies in order to meet demand. Engineering and operational constraints are often included, such as access to and costs of transmission, the availability and quality of renewable resources, ancillary service requirements and their costs, and physical limitations of operating different types of power plants (Short *et al.*, 2009). The detail in these models allows for the explicit and more realistic consideration of how different technologies within the system interact. However, bottom-up models lack an economy-wide

framework and therefore cannot provide measures of economy-wide consumption or policy costs.

3.1.2 Top-Down Economic Models

Top-down models represent economy-wide relationships, can measure social welfare, and are suitable for simulating a wide variety of policies and their impacts. Computable General Equilibrium (CGE) models are a primary type of top-down economic model. CGE models represent the circular flow of goods and services in the economy (Figure 3.1). Consumers (households) supply capital and labor services to the producing sectors, which in turn supply goods and services to consumers. The models also represent the reverse flow of payments that corresponds to the flow of goods and services: households receive payments from the producing sectors for the labor and capital services they provide and in turn use that income to pay producers for the goods and services they consume. CGE models track all of these transactions within and across multiple sectors as well as among different regions. Supply, demand, and prices are determined endogenously by all sectors being in equilibrium and all markets clearing.

(Source: Paltsev et al., 2005)

Figure 3.1 Computable General Equilibrium Circular Flow

CGE models focused on energy and environmental policy exist at various levels of economic aggregation. At a high level of aggregation is the Dynamic Integrated model of Climate and the Economy (DICE) (Nordhaus, 1992), and its extensions such as RICE (Nordhaus & Yang, 1996), ENTICE (Popp, 2004), and ENTICE-BR (Popp, 2006). These macroeconomic

models are built upon the neoclassical Ramsey optimal growth framework, in which growth is driven by capital accumulation and economic equilibrium is reached when the utility function is optimized intertemporally. These models are highly aggregated, often representing the economy with a single sector, or very few, and production of a single final good. Details about the productive inputs (e.g., capital, labor, and energy) are also limited. This level of detail is fit for studying general economy-wide carbon policies such as a global carbon tax or cap and trade system, or as a highly stylized cost-benefit framework. Its transparency is also well-suited for testing representations of difficult concepts, such as technological change. However, the lack of detail does not make it a viable option for studying specific sectors and the choice of technologies within a sector.

Sometimes high level macroeconomic models (similar to DICE) are linked to bottom-up engineering models. Examples of these "hybrid" models are the Model for Evaluating the Regional and Global Effects of GHG reduction policies (MERGE) (Mann *et al.*, 1995b) and the World Induced Technical Change Hybrid Model (WITCH) (Bosetti *et al.*, 2006).

Other economic models follow a more disaggregated, multi-sector economic framework and are constructed from input-out data for the economy. These models are particularly useful for studying the economy-wide impacts of policies as well as sector-specific decisions. Examples are the MIT Emission Prediction and Policy Analysis (EPPA) model (Paltsev *et al.*, 2005), EPA's Applied Dynamic Analysis of the Global Economy (ADAGE) model (Ross, 2008), Purdue University's Global Trade Analysis Project (GTAP) Model (Hertel, 1997), and Charles River Associates' Multi-region National (MRN) model (Smith, 2007).

Overall, CGE models are very powerful tools for assessing the economy-wide impacts of policies because they capture feedbacks throughout the economy. For this reason, it is a particularly appropriate tool to study the impacts of electricity sector strategies and emissions reductions policies. Changes in the electricity sector affect other sectors throughout the economy, and a CGE model can capture those effects. For example, if a policy causes electricity prices to increase, the prices of goods produced using electricity can increase, and consumers may have less money to spend in other sectors. Electricity may also become important to the transportation sector through plug-in electric vehicles, or affect the agricultural sector by using biomass for generation and competing for land resources. The key point is that a CGE model captures all of

these ripple and feedback effects throughout an economy and can therefore provide an accurate estimate of the full economy-wide cost of a policy or strategy.

While focusing on economic details of market flows, CGE models often make simplifications when it comes to technical detail. Common simplifications are to aggregate sectors, include a subset of representative technologies, and make assumptions about the general impact of details that are not explicitly represented. In the electricity sector, CGE models often lack the full suite of technology options as well as operational constraints such as ramping or transmission congestion. Such details are only implicitly included in the relative costs of the represented technologies. However, given the model elements required for the research question in this work, CGE models are a suitable decision-support tool.

3.2 CGE Modeling Approach

For questions about electricity investments and emission reductions in the next ten or so years, a CGE model is not only appropriate, but has several advantages. While explicit inclusion of technical details is crucial to questions about electricity generation on the seconds to daily timescales, the longer timescales of electricity capacity expansion make a CGE model a suitable choice. An advantage of using a CGE model for such questions is that supply, demand and prices are all endogenous instead of exogenous assumptions. More importantly, a CGE model captures the macroeconomics effects and economy-wide welfare costs of electricity investment and emission reduction decisions.

3.2.1 Background

The CGE model has its origins principally in neoclassical modeling developments and invokes microeconomic principles (Arrow & Debreu, 1954; Shoven & Whalley, 1984). Based on their endowments[4] and preferences, one or more representative agents maximize utility subject to a budget constraint, while producers maximize profits with production functions specified as constant returns-to-scale. The equilibrium solution is comprised of a vector of prices and quantities for which demand equals supply, household income equals expenditures, and the

---

[4] In economics endowments refer to sources of income with which an agent is endowed (i.e. owns). Endowments include capital, labor, natural resources, and anything else an agent owns and can earn income on.

profits of firms are driven to zero. The critical data that determine the structure of a CGE model are contained in Social Accounting Matrices (SAMs), which represent a snapshot of the economy of a region in a single model benchmark year. SAMs are developed from National Income and Product Account data and the input-output tables that quantify the inter-industry flows of goods and services (Pyatt & Round, 1985). To study climate policy, CGE models must also represent the physical flows of carbon-based fuels and resources in the economy and their GHG emissions. Many CGE models are written in the GAMS software system and may be formulated in the MPSGE programming language (Rutherford, 1999).

Production functions for each sector describe the ways in which capital, labor, energy and intermediate inputs can be used to produce output. Consumption is modeled as if there were a representative consumer(s) maximizing utility by choosing among goods and services. Typically in CGE models production sectors and final consumption are modeled using nested Constant Elasticity of Substitution (CES) production functions. A fundamental feature in CGE models is the representation of the ability of individuals to make tradeoffs among the inputs to both production and consumption, particularly in response to changes in input prices. For producers this reflects the underlying technology—the extent to which labor, capital and energy can be substituted for each other. The technical ability or willingness to make such tradeoffs is summarized by elasticities of substitution, which are key parameters in production and utility functions. Elasticities of substitution are generally important determinants of the estimates of the cost of policies. For example, if an emissions cap increases the price of carbon-based fuels, the cost of production of goods will rise depending on the energy share, its carbon content and the ability of the industry to substitute other inputs for energy. High elasticities mean producers can easily switch away from the higher cost fuel input, therefore reducing the cost burden of the policy. Similarly, consumers' ability to shift away from the use of energy and energy intensive goods and therefore be less affected by the policy depends on the elasticity of substitution and how prices for each of the goods changes. Elasticity values are typically based on econometric evidence or other methods as appropriate (Arndt *et al.*, 2007; Balistreri *et al.*, 2003; Zhang & Verikios, 2006).

Another important aspect of CGE models is the degree to which they capture the dynamics of the economy over time. Savings and investments, capital accumulation, technological change, resource depletion, and labor productivity growth are key processes that

govern the evolution of the economy and energy use over time. There are different approaches to representing these processes. Particularly important for this work is capital stock accounting. Most CGE models use a putty-putty representation (Phelps, 1963) or a putty-clay capital vintaging structure (Paltsev *et al.*, 2005), depending on whether investment is considered malleable or locked into vintage capital. To reflect the long lifetimes of electricity investments, capital vintaging is required. Capital vintaging for electricity locks investments into place and tracks the amount of electricity generation capacity available from previous years. In this way, today's decisions about how much of each technology to build affect the electricity system long into the future.

CGE models facilitate the computation of measures of the total cost of policies. These take into account multiple feedbacks on production, income and demand across the full range of industries in an economy. One such measure, common in economic analysis, is the change in economic welfare (consumption) measured as equivalent variation. Conceptually, this is the amount of income needed to compensate the representative agent for welfare losses suffered as a result of a policy. Other common outputs of CGE models are commodity prices (oil, coal, natural gas, electricity, etc.), factor prices (wages, capital returns), resource rents (rents associated with oil, gas, and coal deposits and land), the emissions price corresponding to an emissions cap, and the optimal energy mix and electricity mix. Although the optimal electricity mix is solved as if from the perspective of a central planner, one can think of it as the aggregate result of individual firms optimizing their own profits. For a more detailed explanation of how CGE models are formulated and used to study economy-wide policy impacts, see Sue Wing 2004.

Due to their ability to capture economy-wide policy costs, CGE models are well-suited for studying energy and climate policies. At the global level, CGE models can be used to assess scenarios to stabilize GHG concentrations (e.g. Clarke *et al.*, 2007). CGE models are also used to study regional policies, such as the European Union Emission Trading Scheme (e.g. Klepper & Peterson, 2004; Böhringer *et al.*, 2006). Much CGE analysis has been conducted on U.S. policy. Several studies analyzed congressional proposals (e.g. Rausch *et al.*, 2010; EPA, 2009b; EIA, 2009; Paltsev *et al.*, 2009; Metcalf *et al.*, 2008; Paltsev *et al.*, 2008) Electricity-specific policies are also studied using CGE models. For example, some studies look at how a cap-and-trade system covering the electricity sector interacts with a renewable portfolio standard (e.g. Morris *et al.*, 2010; Böhringer *et al.*, 2009) or the distributional and efficiency impacts of an RPS (e.g.

Rausch & Mowers, 2012). Other CGE studies look at the viability and impacts of different electricity generation technologies, such as carbon capture and storage (CCS) (e.g. McFarland & Herzog, 2006), natural gas (e.g. Paltsev *et al.*, 2011), or renewables (e.g. Böhringer & Loschel, 2006). Regardless of the focus of the study, many CGE analyses participate in technology forecasting by reporting the resulting electricity mix as an output of the model. Accordingly, we can see how the optimal the electricity mix differs under different policies, different assumptions and different models. Thus there is a well-established precedent and value of using CGE models to study energy and climate policies and their impacts.

3.2.2 Key Limitations

Despite their many benefits and usefulness, current CGE models also have some key limitations and challenges to overcome in order to improve their relevancy as real-world decision aids.

Uncertainty

According to the U.S. Department of Energy (DOE) Office of Science, one of the main challenges is the representation of uncertainty and risk (Janetos *et al.*, 2009). CGE models are deterministic and therefore do not explicitly capture decision-making under uncertainty. Chapter 2 discussed two critical uncertainties that affect near-term electricity investment decisions: technology costs and government policy. While the supply and demand of electricity and the electricity price are determined endogenously in a CGE model, they are driven by deterministic assumptions. Assumptions about the policies in place, the technologies available and their relative costs, elasticities of substitution in the production and utility functions, technological change, resource availability, labor productivity growth, capital stock accounting (e.g. vintaging and depletion), interactions with other countries (e.g. in terms of trade), and other model parameters have important impacts on the results of a CGE model. Although scenario and sensitivity analysis is often performed, CGE models do not formally include decision-making that considered these uncertainties.

As discussed in Chapter 2, of the many uncertainties involved in electricity and emission reduction decisions, uncertainty in the cost of technologies is among the most important because these costs determine the energy mix, which is a main determinant of emissions and policy costs. The energy sources that could potentially play a significant role in a carbon-constrained future

include wind, solar, carbon capture and storage for coal- or gas-fired power plants, advanced nuclear technologies, and biomass. Each of these technologies has significant uncertainty about future cost, performance, and potential barriers to deployment at scale. However, most CGE analyses consist of deterministic forecasts of future technology mixes, sometimes supplemented with alternative scenarios or sensitivity tests (e.g. Paltsev *et al.* 2009; Clarke *et al.*, 2007; Fisher *et al.*, 2007). A deterministic CGE can tell us the optimal electricity mix under each of the technology cost scenarios, but those mixes are different for each scenario so we can only implement the mix if we know for certain which scenario reflects the real world. If we do not know which scenario reflects the real world and are uncertain about the relative costs of the technologies, the model cannot tell us which mix is best.
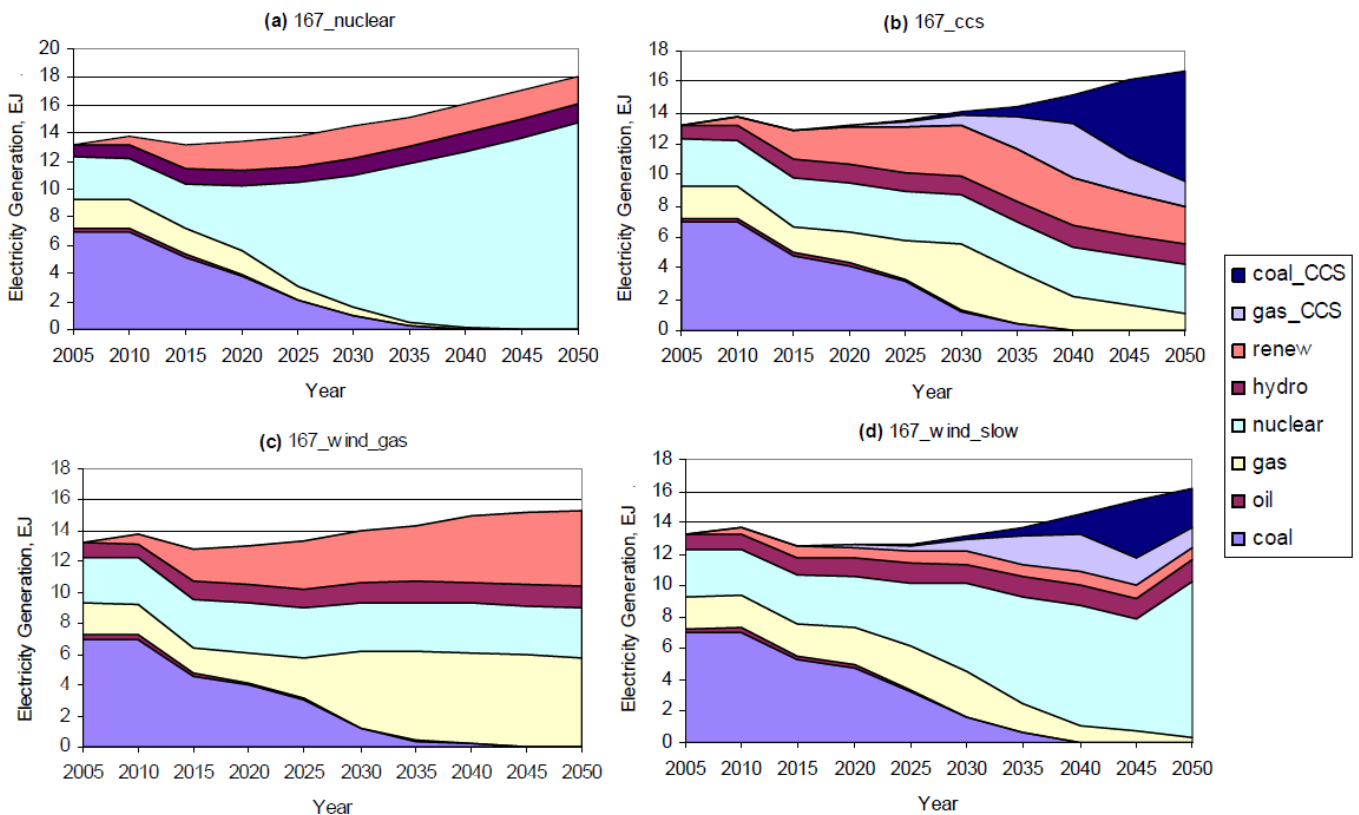
(Source: Paltsev et al., 2009)

Figure 3.2 Alternative Technology Assumptions and Generation Choices

As an example, Paltsev et al (2009) explore a policy reducing emissions by 80% by 2050 under different scenarios changing the relative costs of advanced low-carbon technologies and show that fairly small changes in the relative costs can lead to a very different set of generation

51

choices (Figure 3.2). These scenarios lead to changes in the 2030 economy-wide policy cost of 3-9%, moderated by the choice of other advanced technologies when one or more are made more expensive. The impact on policy cost would be even greater had the study considered scenarios in which all advanced technologies are expensive or unavailable. As this example indicates, the costs of technologies have a large impact on both the desired electricity mix and the cost of meeting policy, suggesting the value of formally incorporating this uncertainty into a CGE model.

Expectations

A related limitation of CGE models is how the expectations of modeled decision makers are represented. The majority of CGE models make one of two assumptions. Some models assume forward-looking behavior with perfect foresight, in which decisions over the entire time horizon are optimized at once. Examples of this type of model include MERGE (Manne *et al.*, 1995b) and DICE (Nordhaus, 1992). Other models are myopic recursive-dynamic, in which decisions are made in each time period sequentially without information on what will happen in future periods. The MIT EPPA model (Paltsev *et al.*, 2005) and MiniCAM (Edmonds & Reilly, 1985) are examples of this type of model. Neither of these representations of expectations reflects the real world. In reality we have uncertain expectations about the future that we update as we learn new information.

Policy uncertainty clearly raises the question about what should be done in the years before the policy starts. Does the decision maker fully anticipate the policy and plan accordingly (forward-looking) or completely ignore that a policy may be in place in the future (myopic recursive-dynamic). In a forward-looking CGE model with capital vintaging the electricity mix chosen in years before the policy starts will reflect the knowledge that the policy is coming.[5] In a myopic recursive-dynamic model the mix chosen in years before the policy starts will not reflect any assumptions that a policy is coming (i.e. they will continue to reflect a "business as usual" world).

---

[5] Most existing forward-looking CGE models either do not have capital vintaging or vintage current capital but assume all future capital is malleable (not locked in place). These models focus on consumption smoothing— changing savings and investment decisions in anticipation of higher consumption costs in the future due to policy. Without capital vintaging (i.e. a putty-putty structure instead), there would be no reason to change the electricity mix in anticipation of future policy because it can be costlessly reallocated when the policy begins.

A more realistic alternative to either of these expectation frameworks, is one in which the decision maker has probabilistic expectations about future conditions and makes decisions while directly taking the uncertainties into account. So in the case of policy uncertainty, instead of running one policy scenario at a time, we run a scenario of a probability distribution of potential policies (e.g. 50% chance of no policy, 30% chance of a 20% cap, and 20% chance of a 40% cap). In years before the policy starts the decision maker sees those probabilities about what policies might be in place in future years and makes decisions about what the electricity mix should be taking that into account. The decision maker does not operate on perfect information or no information, but on imperfect information and hedges decisions accordingly. The next section discusses methods specifically designed to represent this type of sequential decision making under uncertainty.

3.3 Decision Making Under Uncertainty

3.3.1 Overview of Methods

Many analyses tend to ignore uncertainty and apply deterministic models, often using the expected value of uncertain parameters. A common approach is to perform sensitivity or scenario analysis, in which one assumes different values of parameters and re-solves the deterministic model for each. Monte Carlo simulation is another option in which repeated random draws from probability distributions of parameter values are made and used to re-solve the deterministic model. While these methods can be useful for some questions, they do not represent *decision making under uncertainty*. In most real-world situations we do not know the correct value of a parameter but have an idea of its probability distribution and over time we can learn information about the true value. In the meantime, decisions are made taking the uncertainty into account.

There are several methods that explicitly frame problems as sequential decisions under uncertainty. The main approaches are dynamic programming and stochastic programming. Both take use probability to represent uncertainty, but differ significantly in their implementation. These methods however become computationally intractable for problems beyond a certain size. Bellman coined this problem "The Curse of Dimensionality" (Bellman, 1957). Newer methods of approximate dynamic programming (ADP) and stochastic dual dynamic programming (SDDP) seek to overcome this problem. ADP and SDDP combine regular dynamic programming

and stochastic programming approaches, respectively, with Monte Carlo sampling techniques. A survey of these methods and their strength and weaknesses is useful to understanding the overall landscape of decision making under uncertainty.

Dynamic Programming

Dynamic programming (DP) is an approach to modeling and solving dynamic and stochastic decision problems, originally developed by Richard Bellman (Bellman, 1957). Dynamic programs are also referred to as Markov decision processes (MDP) when decision periods are discrete (Puterman, 2005) and stochastic optimal control problems when decision periods are continuous (Bertsekas, 2007). Although DP can be useful in solving large deterministic dynamic problems, often DP actually refers to stochastic dynamic programming (SDP) because it is used to model problems under uncertainty. The structure of DP problems focuses on the state of the system and the value function, and assumes a Markov process. The Markov property is "memoryless" in that future value depends solely on the current state of the system, not its full history. DP is a decomposition method for the optimization of large, non-convex problems, in which the problem is broken down into a sequence of simpler subproblems: now and later. The main reason why it is important to distinguish between the decisions made at different points in time, is that the information available to the decision maker changes over time—typically more information is available at later decision points. DP models can be divided into Finite Horizon Problems, in which optimal strategies and value functions depend on time, and Infinite Horizon Problems, in which the steady-state optimal strategy and value function do not depend on time. Finite horizon problems are appropriate for problems of transitions or where there is a specific terminal period (e.g., a baseball game or a chess game), and are solved by backward induction (more detail below). Infinite horizon problems are used when the stochastic processes and state-transition processes are stationary and where there is no specific ending (e.g., inventory policy for a business on any given week). Infinite horizon problems can be solved using a variety of methods, including linear programming approaches, value iteration, policy iteration, or a hybrid of value and policy iteration.

A fundamental element of DP models is the value function, also called the objective or utility function. The value function is a way to value and compare different outcomes of decisions and uncertainty realizations. In DP problems the goal is typically to maximize the

expected sum of the rewards obtained in each time period, or equivalently, to minimize the expected sum of the costs incurred in each time period. The Bellman equation (Bellman, 2003) describes the optimality conditions that must hold simultaneously for all decision stages and for all states: in each time period you choose an action that maximizes the current reward plus the expected discounted future reward:

$$V_t(S_t) = \max_{x_t} \left[C_t(S_t, x_t) + \gamma E\{V_{t+1}(S_{t+1}(S_t, x_t) | S_t\}\right] \qquad \text{(EQ. 3.1)}$$

where:
t is decision period,
V is total value,
S is state (summary of available information that affects stochastic process, e.g. electricity capacity level at time t),
C is consumption,
x is decision (or action),
γ is discount factor.

Equation 3.1 states that for each time period t, the optimal value of being in a given state is determined by choosing the action x_t that maximizes current consumption C_t (which is a function of the current state S_t and the action chosen now) *plus* the discounted expected optimal value in the next time period V_{t+1}, which is a function of the expected next state S_{t+1} given the current state and the action chosen in the current time period.

This finite horizon problem is solved by backward induction. To do this the modeler must know or assume the terminal conditions—the value of being in each possible state in period T (where T is the final year of the horizon). Starting in T-1 and using backward induction, you compare the values of being in each state, taking expected values of states branching off of the same uncertainty, and fold back to the previous time period to choose the action that will put you in the best position to achieve the highest expected value in the next period. You continue folding back until you arrive at period one at which point you choose the action that best optimizes the expected value of all future time periods.

Overall DP is a very useful method for decision making under uncertainty and capturing the dynamic and stochastic nature of many real-world problems. Decisions are made with the information currently available and expectations about the future which are updated as we learn over time. DP can handle a large range of problem types, including those that are non-convex. It allows for flexibility in the setup, including the ability to allow decisions to affect the probability

distributions of uncertain variables. Although the Markov property is assumed, path dependency can be captured by augmenting the state variables to include past information. The main challenge to the application of DP is its potential to run into the curse of dimensionality. It requires that all possible states, actions and uncertainties are enumerated for each time period. In terms of a decision tree, each possible path of the tree must be solved. The problem explodes exponentially as the state, action, uncertainty, and time space is increased, and problems can quickly grow to become computationally infeasible. However, with advances in computing power, DP is still able to represent realistic problems of considerable size.

Stochastic Programming

Stochastic programming (SP) is another approach for modeling optimization problems under uncertainty ((Birgre & Louveaux, 1997; Kall & Wallace, 1994). Stochastic programming builds on linear programming methods, although there are variants for dealing with non-linear and integer problems. Because SP takes advantage of optimization and mathematical programming techniques, convexity is a key requirement. Convexity ensures that any local minimum is a global minimum. In SP uncertainty is modeled as a random variable with a probability distribution that is known or can be estimated. In order to maintain convexity, actions cannot affect probability distributions. The goal is often to find some policy that is feasible for all (or almost all) the possible parameter realizations and optimizes the expectation of a given objective function of the decisions and the random variables The most commonly studied SPs are two-stage stochastic programs with recourse. The optimal policy from such a model is a single first-stage decision and a collection of recourse decisions (a decision rule) defining which second-stage action should be taken in response to each random outcome. This setup can also be extended to multiple stages.

SP can be formulated and solved using several approaches. The deterministic equivalent approach optimizes the expected value by solving simultaneously for all scenarios, and includes non-anticipativity constraints for first-stage decisions. A typical implementation of this approach indexes variables not only by time, but also by state of the world (SOW). This is a very useful approach which is easy to implement, but it can only handle smaller sized problems, because the size of the linear program being solved increases with the number of scenarios and stages. SPs can also be solved by one of several decomposition techniques which break the problem up into smaller problems. Benders Decomposition (or the "L-shaped" method) is a widely used

technique. It uses an exogenous scenario tree, and formulates two subproblems, the master problem for the first-stage decision and the slave or subproblem for the second stage problem for each scenario. The algorithm iterates between the two problems, adding new constraints ("Benders cuts") until it converges to an optimal solution (Birgre & Louveaux, 1997). With this formulation, the number of variables can be reduced substantially and subproblems can be solved in parallel. However, like other approaches discussed, this approach also blows up as dimensions increase and for multiple stages. Other SP solution approaches include Dantzig-Wolfe (Dantzig & Wolfe, 1960), Lagrangian Relaxation (Lemaréchal, 2001), and Progressive Hedging (Rockafellar & Wets, 1991).

Overall SP is a very useful method for decision making under uncertainty. It takes advantage of linear structures, convex analysis and duality theory for efficient optimization. The tools of mathematical programming are also indispensible in handling general constraints on states and decision variables. The addition of constraints is often a serious impediment to dynamic programming techniques as it increases the dimension of the state space, which can lead to an intractable problem. SP generally allows for continuous actions, but discretized uncertainties. Path dependency is natural to capture in an SP framework. However, the need for convexity, the issue of convergence, and the difficultly of representing action-dependent probabilities are limitations. Also, as for DP, the main problem with SP is the curse of dimensionality. The problem size grows exponentially as the scenarios (action and uncertainty space) and time dimensions are increased.

Approximate Dynamic Programming

Approximate Dynamic Programming (ADP) is a class of methods developed to address the Curse of Dimensionality (Powell, 2011; Bertsekas & Tsitsiklis, 1996). ADP is also sometimes referred to as neuro-dynamic programming, forward dynamic programming and adaptive dynamic programming. ADP combines traditional DP with Monte Carlo sampling and response surface approximation strategies. The idea is that instead of exhaustively searching through all possible states, decisions, and information signals, ADP samples possible paths through a scenario tree to construct an approximation of the value function which can then be used to make optimal decisions for any possible state. ADP methods emerge from operations research. Recent work (e.g., Webster *et al.*, 2012a; Godfrey & Powell, 2002; Basler, 2006;

Powell *et al.,* 2012) has successfully implemented ADP methods on large-scale, multi-dimensional problems.

There are many different variations on ADP. The representation of the value function approximation is one of the critical design choices. Common approaches for value function approximation are recursive linear regression (e.g., Ormoneit & Sen, 2002), piecewise-linear (e.g., CAVE, Godfrey and Powell, 2002), moving least squares or mesh-free (Parpas & Webster, 2011), neural networks (e.g., Bertsekas & Tsitsiklis, 1996), and aggregation of states (e.g., George & Powell, 2006). Sampling strategies are also key to ADP. The goal is to sample states in order to minimize sampling states that give no additional useful information while making sure to sample states never visited that would give information. This is known as the "explore vs. exploit" problem (Powell, 2011). The algorithm structure can also take different forms. Single-pass procedures (iterate forward through time periods for each sample), double-pass procedures (iterate forward through time to find decisions, then backwards to compute value), temporal difference learning (hybrid between single- and double-pass), policy iteration, and value iteration procedures are all options.

The main advantage of ADP is that it overcomes the curse of dimensionality. Also, because it builds on DP methods, it can handle a large range of problem types and decisions can affect probabilities. ADP is efficient because it tries to find promising space and policies for that space without finding the whole value function. Of course, the corresponding drawback is that it may not converge to the true answer because it only approximates a global value function. Thus convergence becomes an issue. Also the only way to get statistics or bounds for the solution is to run the ADP model multiple times, which offsets some of its efficiency. Also, there may be a need to go through every possible action to find an optimal when building the value function approximation because it is not a linear program. On the upside, this allows the handling of non-convex problems because going through all actions guarantees a global optimum. There are techniques that can alleviate the need of looping through all actions. There are also many different solution options for ADP, though they may be problem dependent.

Stochastic Dual Dynamic Programming

Stochastic Dual Dynamic Programming (SDDP) is another promising method that can overcome the Curse of Dimensionality. SDDP combines traditional SP with Monte Carlo sampling. SDDP is essentially multistage Benders with sampling. Regular SP with Benders

requires that you explicitly build a scenario tree with exogenous scenarios which is quite costly computationally and infeasible when the problem gets too big. The key to SDDP is that it gets around building a tree by sampling. Scenario paths are samples and then Benders cuts are performed on the samples to approximate the value function. In this way SDDP can overcome the curse of dimensionality. Also, because it builds on SP methods and linear programming, SDDP can take advantage of mathematical programming techniques and clever optimization techniques (because it always deals with linear formulations of problems). Of course the corresponding drawback is that it requires that problems are convex, and actions cannot affect probability distributions. Because SDDP is approximating the expected value of the value function, it computes an exact global solution, which enables the calculation of statistics and bounds. Also, convergence can be a challenge.

Summary

There are several methods for approaching problems of decision making under uncertainty. Each has its pros and cons. The main challenge for traditional DP and SP approaches is the curse of dimensionality. ADP and SDDP are ways to overcome that curse. It is interesting that the distinct camps of DP and SP are beginning to merge back together, driven by the need to combat real-world large-scale, multi-dimensional problems. ADP and SDDP apply essentially the same algorithm at a general level; both focus on sampling and approximating a value function. The best method to use depends on both the research question and the model being used.

For the work in this dissertation, which employs a small CGE model to study near-term electricity and emission decisions under uncertainty in policy and technology costs, a DP approach is ideal. DP is well-suited for the non-linear, non-convex nature of the energy-climate problem and the CGE model is small enough to explore a wide range of decisions and uncertainties without running into dimensionality problems. A full description of the model and DP approach used in this work is provided in the next chapter.

3.4 Applications of Uncertainty Methods to CGE Models and Limitations

There have been numerous applications of uncertainty analysis methods to engineering-cost models in the study of electricity capacity expansion. Sensitivity analysis, scenario analysis and Monte Carlo simulations are common approaches (e.g. Bergerson & Lave, 2007; Richels & Blanford, 2008; Blanford 2009; Messner *et al.*, 1995; Grubb, 2002; Grubler & Gritsevskii, 1997). Formal DP and SP approaches have also been applied to engineering-cost models (e.g. Klein *et al.*, 2008; Mattsson, 2002; Ybema *et al.*, 1998; Bosetti & Tavoni, 2009; Kypreos & Barreto, 2000; Botterud *et al.*, 2005). There have even been a few applications of ADP to engineering-cost models (e.g. Santen, 2012; Powell *et al.*, 2012). While all of these studies offer valuable insight and are appropriate for certain research questions, their bottom-up nature prevents them from addressing questions of economy-wide impacts of policies and investment decisions. This research seeks to study decision-making under uncertainty in which decisions are based on societal economy-wide costs, not just direct costs to producers in a specific sector. To do that, a CGE model must be employed.

Although significant research has gone into the development of CGE models and their application to policy analysis, less work has focused on the representation of decision-making under uncertainty with learning in CGE models. Notable work has applied CGE models to probabilistic uncertainty analysis, typically in the form of Monte Carlo simulation (e.g. Webster *et al*., 2009; Webster *et al*., 2008a; Scott *et al*., 1999; Manne and Richels, 1994; Reilly *et al*., 1987). This type of analysis is quite useful in that it provides probability distributions of climate impacts and economic costs of policies. However, in this approach each scenario is still modeled as though the randomly sampled parameter values are known and correct, and decision-making under uncertainty with learning is not captured.

There have been some studies using CGE models that explicitly frame energy or emission decisions as sequential decisions under uncertainty, thereby capturing realistic expectations about the future and the ability to adjust decisions as information becomes available. These studies employ dynamic programming or stochastic programming techniques. However, application of these approaches to CGE models has been extremely limited and the curse of dimensionality has required that analyses reduce the size of the problem in some way. There are two main ways of doing so: reduce the dimensions or use a simplified model.

There are several ways to reduce the dimensions of the problem. Discretization can reduce problem size by limiting the values the state variables can take on, the actions that can be taken (for example, electricity can only be generated in blocks of 100 kWh or 200 kWh), or the uncertain parameter values. One can also eliminate from the exploration space state values that are known to be infeasible (for example, it is impractical that renewable sources could increase to 50% of total generation in the next 5 years, so those scenarios can be omitted). One can also limit the number of possible decisions. For example, instead of choosing amounts for many technologies, can be limited to two technologies—one fossil-based and one a non-fossil. Representing fewer possible uncertainty realizations is another option, for example high, medium and low technology cost. Representing a limited number of uncertain states-of-the-world (SOWs) is an approach that has been applied to CGE models for energy/climate policy analysis. (e.g. Webster, 2008c; Webster, 2002). Another approach is to reduce the time dimension by limiting the number of decision periods. Decisions could be made less often, for example every 5 years, 10 years, 20 years, etc. Two-stage representations are also common. In stage one, decisions are made under uncertainty and in stage two the uncertain outcome becomes known and a second decision is made (e.g. two-stage decision trees, or two-stage stochastic programming with recourse). Manne and Richels (1995a) and Hammitt *et al.* (1992) utilize the two-stage "act-then-learn" framework.

The main advantage of methods that reduce or discretize the dimensions to reduce the problem size is that it can allow the problem to be solved using traditional dynamic or stochastic programming methods. Prior to recent methodological advances (e.g. ADP and SDDP), this was the only option available for solving high-dimensional problems. Basically the problem had to be reduced to a size for which the curse of dimensionality is not a problem and traditional DP or SP methods could be used. The advantage of DP and SP methods is that they are well-established, and relatively easy to use. The obvious disadvantage of reducing the dimensionality is that the simplifications can prevent a full exploration of decisions under uncertainty. Low resolution discretization of states, actions and uncertainties introduce artificial non-linearities into the results. Accordingly, the optimal policy approximated from this approach may be a poor approximation since the full range of uncertainties, decisions, and states was not explored. Alternatively, depending on the discretization, the solution may not even be implementable. For example, the optimal period 1 decision could be in ranges, such as generate 10-15 EJ electricity

from coal, 5-10 EJ from gas, 0-5 EJ from wind, etc. However, if reasonable limitations are made, this approach can still provide satisfying resolution.

The other main approach to contain the size of the problem is to use a simplified or stylized model. The advantage of this approach is that it may allow for a DP or SP solution of an optimal policy without reducing the dimensionality of the state, decision or uncertainty space. A smaller model is more tractable to solve for larger numbers of dimensions. Another advantage is that the model may be more transparent, making it easier to identify the sources of results. Because the DICE model is so aggregated, it is more amenable to DP and SP formulations. For this reason DICE has been widely used to study questions of uncertainty (e.g. Gerst *et al.*, 2010; Kelly & Kolstad, 1999; Leach, 2007; Crost & Traeger, 2010; Lemoine & Traeger, 2011; Nordhaus & Popp, 1997; Kolstad, 1996; Yohe *et al.*, 2004; Webster *et al.*, 2008b). The drawback to this approach is that a simplified model may lack a satisfactory level of detail. However, if there is sufficient detail in the model areas directly relevant to the question of interest, this approach can be a sensible choice.

Recent methodological advances (e.g. ADP and SDDP), present the prospect of allowing the explicit representation of uncertainty in a CGE model without limiting the size of the problem. Because the electricity and climate policy problem is non-linear and not necessarily convex, it would be difficult to fit into an SDDP framework. While it may be possible to extend SDDP methods to problems of this type (just as regular SP has been extended to non-linear and integer problems), the method has not been developed to that point yet. ADP is more flexible than SDDP in that it can handle a large range of problem types (non-linear, non-convex, etc.) and decisions can affect probabilities. It holds great potential for solving a high dimensional energy and climate decision problem over a long time frame with multiple uncertainties, multiple decisions, the ability to learn, and a model that capture rich details and economic interactions (e.g., Godfrey & Powell, 2002; Basler, 2006; Bertsekas, 2007; Parpas & Webster, 2011; Powell *et al.* 2012). Webster *et al.* (2012a) is the only known work applying ADP methods to a CGE model (DICE). This work incorporates endogenous, path-dependent continuous uncertainty in technology change in a multi-stage decision context. It shows the value of the ADP framework to study optimal decisions under uncertainties and to explicitly overcome the dimensionality challenge of SDP without loss of accuracy. Despite limited applications, there is great potential for application to CGE models of greater resolution than DICE.

This dissertation contributes to the literature by applying the stochastic dynamic programming method to a small CGE model to investigate near-term electricity investment and emission reduction decisions under uncertainty in policy and technology costs. In using a simple and transparent model which still captures important details, it seeks to demonstrate the ability to and importance of capturing decision making under uncertainty in a CGE model to investigate problems such as the one in this dissertation. Given the size of the model used, it is amenable to a DP approach that maintains a high resolution of the decision space. This work represents the first time DP and CGE modeling approaches have been merged to study this research question. The next chapter introduces this new DP-CGE modeling framework.

(This page is intentionally left blank)

Chapter 4: Modeling and Analysis Framework

To study the research questions posed above, a small-scale CGE model has been formulated as a stochastic dynamic program (DP). The DP-CGE model was built with Professor John Little's advice in mind that models should be as simple and transparent as possible while also being complete on important dynamics (Little, 1970). A small, transparent model is particularly valuable for the purposes of demonstrating the role of uncertainty in decision-making.

The model focuses on a single region representing the U.S. It contains a small number of sectors that are important to energy, including electricity and natural resources, and aggregates the rest of the economy into one combined sector. The focus is the electric power sector, which is divided into conventional and low-carbon generation. Within conventional generation there is a choice between coal and natural gas. This stylized model highlights the tradeoffs most relevant to power system decisions under climate policies: conventional vs. low-carbon generation, coal vs. natural gas generation, and new capacity investments vs. changing the operation of existing capacity. The DP formulation allows study of the additional tradeoff of now vs. later in capacity investments and emissions reductions decisions. This chapter first reviews the equilibrium structure of CGE models. Section 4.2 then describes the base CGE model developed for this dissertation, which is deterministic with no foresight. The chapter then describes how the model is reformulated as a DP in Section 4.3. Section 4.4 presents the characterization of the uncertainties and how they are implemented in the model. Finally, Section 4.5 presents details of the model implementation.

4.1 Equilibrium Structure in CGE Modeling

As discussed in Chapter 3, CGE models are based on neoclassic microeconomic principles and represent the circular flow of goods and services in the economy. They track all of these transactions within and across the sectors and regions represented. Supply, demand, and prices are determined endogenously. The critical data that determine the structure of a CGE model are contained in Social Accounting Matrices (SAMs), which represent a snapshot of the economy of a region in a single model benchmark year. SAMs are developed from National Income and Product Account data and the input-output tables that quantify the inter-industry

flows of goods and services (Pyatt & Round, 1985). To study climate policy, CGE models must also represent the physical flows of carbon-based fuels and resources in the economy and their GHG emissions. Many CGE models are written in the GAMS software system and may be formulated in the mathematical programming system for general equilibrium analysis (MPSGE) (Rutherford, 1999).[6]

Equilibrium in a CGE model (as well as in theoretical general equilibrium economics) is achieved when supply equals demand in all markets of the economy in a way that maximizes total economy-wide (consumer plus producer) welfare. This is achieved by firms maximizing profits and the representative consumer maximizing utility.

In each region (indexed by the subscript r) and for each sector (indexed interchangeably by i or j), representative firms choose the level of output y, quantities of primary factors k (indexed by f) and intermediate inputs x from other sectors j to maximize profits subject to the constraint of its production technology:

$$\max_{y_{ri}, x_{rji}, k_{rfi}} \pi = p_{ri} y_{ri} - C_{ri}(p_{ri}, w_{rf}, y_{ri}) \quad s.t. \quad y_{ri} = \varphi_{ri}(x_{rji}, k_{rfi}) \quad \text{(EQ. 4.1)}$$

where:
π = profit function,
C = cost function,
φ = production function
p = price of goods,
w = price of input factors.

In each region, a representative agent is endowed with the supplies of the factors of production k, the services of which may be sold or leased to firms. In each period, the representative agent chooses consumption and saving to maximize a welfare function (W) subject to a budget constraint given by the level of income M:

$$\max_{d_{ri}, s_r} W_{ri}(d_{ri}, s_r) \quad s.t. \quad M_r = \sum_f w_{rf} K_{rf} = p_{rs} s_r + \sum_i p_{ri} d_{ri} \quad \text{(EQ. 4.2)}$$

where:
d = final demand for commodities,
s = saving,
K = aggregate factor endowment of representative agent.

---

[6] MPSGE was developed to support CGE modeling and only requires the user to specify the structure of production and consumption, base data, and parameter values. Based on those inputs the equations for general equilibrium are formulated.

Thus equilibrium is achieved by a representative agent maximizing total economy-wide welfare subject to the budget constraint and technical constraints on how factor inputs can be used in production. Additional constraints are created by policies, such as those limiting emissions. A feasible competitive equilibrium consists of a non-negative consumption plan for the representative consumer, a non-negative production plan for each firm and non-negative prices.

The equations above lead to the following optimality conditions for equilibrium:

(1) *The zero profit condition* requires that any activity operated at a positive output must earn zero profit (*i.e.*, value of inputs must be equal to or greater than value of outputs). This condition means that either output is positive and profit is zero, or profit is negative and there is zero output.

(2) *The market clearance condition* requires that any good with a positive price must have a balance between supply and demand and any good in excess supply must have a zero price.

(3) *The income balance condition* requires that for the representative agent the value of income must equal the value of factor endowments and tax revenue (e.g. from a price on carbon emissions).

The ultimate equilibrium solution is comprised of an endogenously-determined vector of prices and quantities for which demand equals supply (market clearance), household income equals expenditures (income balance), and the profits of firms are driven to zero.

4.2 CGE Model

The CGE model developed in this dissertation follows the structure of the MIT Emissions Prediction and Policy Analysis Model (EPPA) (Paltsev *et al.*, 2005), though considerably simplified. The model is written in General Algebraic Modeling System (GAMS) format and is formulated in MPSGE.

A single region is modeled, approximating the United States in terms of overall size and composition of the economy. The production and household sectors represented are shown in Table 4.1. There is a single representative consumer that makes decisions about household consumption. There are six production sectors: crude oil, refined oil, coal, natural gas, electricity,

and other. Other, which includes transportation, industry, agriculture, etc., comprises the vast majority of the economy (almost 97%). The factors of production included are capital, labor and natural resources (crude oil, coal, and natural gas). Carbon dioxide (CO_2) emissions are included. The amount of emissions resulting from total use of each fossil fuel is defined in the base year, and the amount of emissions per unit of fuel is assumed constant over time. Coal produces over two times the CO_2 emissions of natural gas. Two electricity generation technologies are represented: conventional and an advanced "low-carbon" technology. Conventional electricity can be generated from either coal or natural gas.[7] The low-carbon technology does not use fossil fuels and is more expensive. The underlying social accounting matrix (SAM) data is based on GTAP 5 (Hertel, 1997; Dimaranan & McDougall, 2002) data recalibrated to approximate 2010, which is used as the base year for the model (Appendix A). The model is calibrated on the SAM data to generate a benchmark equilibrium in 2010.

Table 4.1 Model Details

Sectors	Factors
Non-Energy	Capital
Other	Labor
Energy	Coal Resources
Electric Generation	Natural Gas Resources
Conventional (Coal and Natural Gas)	Crude Oil Resources
Low-Carbon	
Coal	
Natural Gas	
Crude Oil	
Refined Oil	
Household	
Household Consumption	

4.1.2 Production and Consumption

Production functions for each sector describe the ways in which capital, labor, energy and intermediate inputs can be used to produce output, and reflect the underlying technology. The consumer utility function describes the preference for each good and service and how they contribute to utility (welfare). The MPSGE software allows users to represent production and

---

[7] Conventional electricity aggregates all generation in the base year, including nuclear, hydro and other generation.

consumption using Constant Elasticity of Substitution (CES) production functions, which take the general form exemplified in Equation 4.3 for the case with two inputs—capital K and labor L.

$$Q = [aK^r + (1-a)L^r]^{\frac{1}{r}} \qquad \text{(EQ. 4.3)}$$

where:
Q = output,
a = share parameter,
K = capital input,
L = labor input,
$r = \frac{(\sigma-1)}{\sigma}$,
$\sigma = \frac{1}{(1-r)}$ = elasticity of substitution.

In addition to capital and labor, other inputs to production include natural resources and intermediate inputs (output from other sectors). By using CES functions, production technologies exhibit constant returns to scale, which implies that firms make zero profits in equilibrium.

The elasticity of substitution reflects the ability to make tradeoffs among the inputs to both production and consumption in response to changes in input prices. Formally, it is the percentage change in input shares (a_i) due to a change in the marginal products of inputs $\left(\frac{\partial Q}{\partial x_i}\right)$. Special cases of the CES formulation are when the elasticity of substitution is 1 (Cobb-Douglas) or 0 (Leontief, meaning they are required in fixed proportions and there is no substitution). A limit of the CES function when expanded to more than two inputs is that that there is a single sigma and so the elasticity of substitution between all pairs of inputs is the same. To overcome this limit it is common to represent a production function as a series of input "nests". Figure 4.1 provides an example of a CES nesting structure: the Capital-Labor nest is defined by a CES function with elasticity of substitution $\sigma_{K\text{-}L}$ and feeds into the top-level Energy-Capital-Labor nest between the KL aggregate and energy which is defined by a CES function with elasticity of substitution $\sigma_{E\text{-}KL.}$
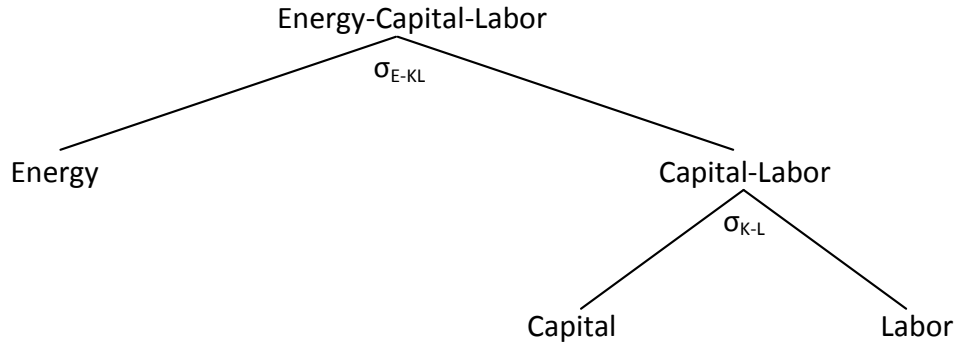
Figure 4.1 Example of Nesting Structure for Constant Elasticity of Substitution (CES) Functions

Elasticities used in this model were guided by those in the MIT EPPA model (Paltsev *et al.*, 2005), which come from review of the literature and expert elicitation. Key elasticities of substitution (σ) are given in Table 4.2. Note that the higher the elasticity, the easier it is to substitute between inputs.

Table 4.2 Elasticities of Substitution

σ	Value	Description
$\sigma_{K\text{-}L}$	1.0	Between capital and labor
$\sigma_{E\text{-}KL}$	0.4	Between energy and capital and labor
σ_{OKL}	1.0	Between other, capital and labor
$\sigma_{ELE\text{-}FE}$	0.5	Between electricity and final energy inputs
σ_{FE}	2.5	Between final energy inputs
$\sigma_{RES\text{-}OTH}$	0.6	Between natural resources and other inputs
σ_{FF}	0.5	Between fixed factor and other inputs
$\sigma_{TOP\text{-}FD}$	0.3	Between aggregate and other goods in final demand
$\sigma_{ENE\text{-}FD}$	0.5	Between energy commodities in final demand

Figures 4.2-4.6 depict the nested structure of the production and consumption functions. A common nest structure exists for extraction of the three fossil fuel resources—crude oil, coal and natural gas (Figure 4.2). The production structure is composed of 3 nests. Starting at the bottom, capital and labor are combined in a CES function to form a capital-labor aggregate, which is then nested with the "other" intermediate input, and this aggregate then appears in a nest with the oil, gas, or coal resource input. The output of these sectors is then the crude oil, gas, or coal produced. In the diagrams all horizontal lines (for example, the Capital-Labor-Other branch in fossil fuel production) represent an elasticity of zero, meaning there is no substation between those inputs.

70
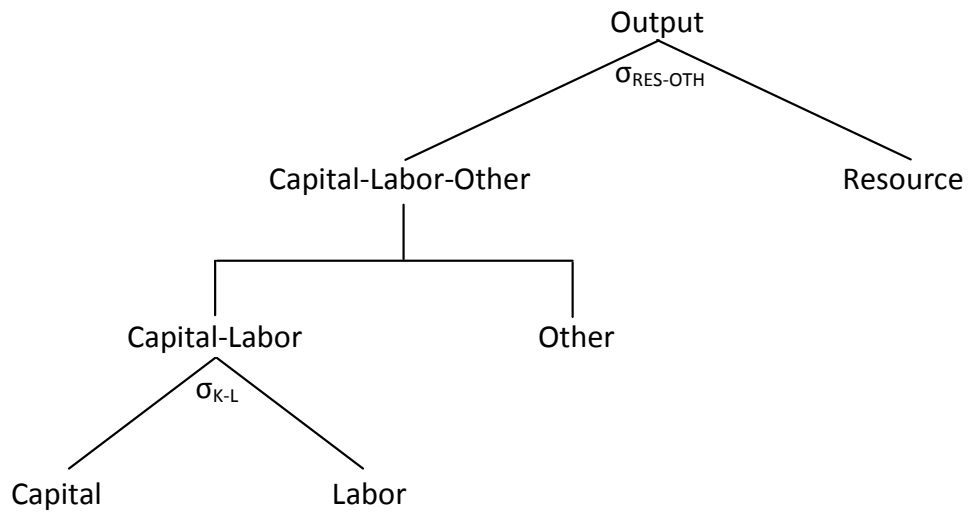
Figure 4.2 Fossil Fuel Production: Crude Oil, Coal and Natural Gas

Figure 4.3 "Other" and Refined Oil Production

The general structure for production for other sectors is shown in Figure 4.3. This includes the "Other" sector, which is the non-energy component of the economy, and the Refined Oil production sector. These sectors include a fuels nest that then substitutes for electricity in an energy nest, which then substitutes for a Capital-Labor aggregate. Crude oil enters along with Other in the top nest with Energy-Capital-Labor in a Leontief structure. This reflects the use of crude oil as a feedstock for the refinery sector. For the "Other" sector the input share of crude oil share is zero and so none is actually used there.

The two electricity producing technologies are conventional and low-carbon generation. There is a single conventional electric technology that uses coal and natural gas as its fuel. Trade-offs between gas and coal generation are represented by the ability to substitute these fuels in generation (Figure 4.4). There is high elasticity of substitution between coal and natural gas inputs ($\sigma_{FE} = 2.5$). This reflects the fact that new investments in electricity can easily switch between coal and natural gas depending on which is cheaper. When an emissions pricing policy is in place, the use of the fossil fuels also requires payment for the carbon emissions that will be produced by the fuel, shown by the Leontief nest of the Carbon Permit and the fuel.

Figure 4.4 Conventional Electricity Production

The low-carbon electricity generation technology (Figure 4.5) produces no carbon emissions and represents advanced low-carbon technologies like wind, solar, CCS, and advanced nuclear. These technologies have little or no market penetration at present, but could take significant market share in the future under some energy price or climate policy conditions. The low-carbon technology is modeled as a single generalized low-carbon technology instead of modeling multiple specific technologies like wind, solar, etc. This representation avoids making assumptions about which low-carbon technology(ies) will be most cost-competitive. Instead it highlights the importance of the relative costs of conventional and low-carbon technologies. Investment in low-carbon generation reflects whichever technologies end up being least costly. The electricity produced from the generalized low-carbon technology is a perfect substitute for conventional electricity (they produce an identical output). The low-carbon technology is not used in the base year and only enters the market if and when it becomes economically competitive. The low-carbon technology initially (in the base year) has a higher cost than conventional generation. This initial higher cost is set in the model by a markup, which is the cost relative to the conventional generation against which it competes in the base year. All inputs are multiplied by the markup. The markup is set to 1.5, indicating that the low-carbon technology is 50% more expensive than conventional electricity. As the prices of inputs change over time, so too does the relative cost of the technologies. The low-carbon technology also requires an additional input of a fixed factor resource, which is described below.
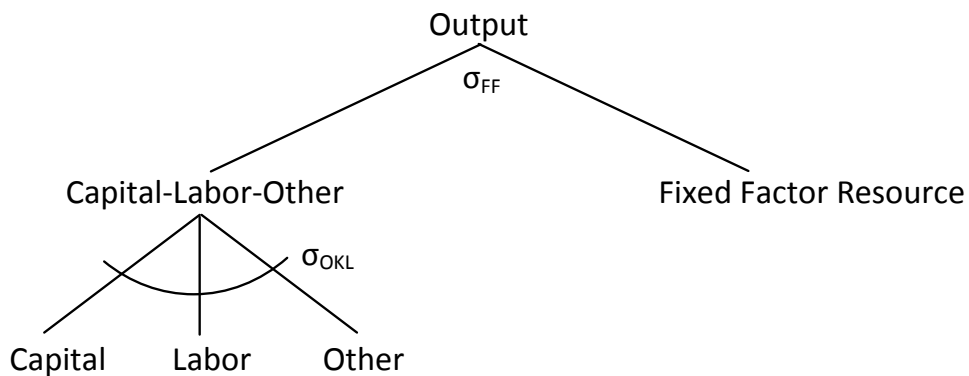
Figure 4.5 Low-Carbon Electricity Production

Figure 4.6 shows the consumer utility function. Investment enters directly into the utility function, which generates demand for investment and makes the consumption-investment decision endogenous. In this model savings are equal to investments. Change in aggregate consumption is an equivalent variation measure of welfare in each period. Households (representative consumer) own all factors of production (including fossil fuel resources), which they "rent" out to producers at a rental rate (the market price for the factors). The total quantity of each primary factor the households are endowed with is based on the data set. That amount combined with the prices of the factors determines the total income for the household. Households then make consumption and investment decisions to maximize their utility preferences subject to their income (as depicted in Equation 4.2).
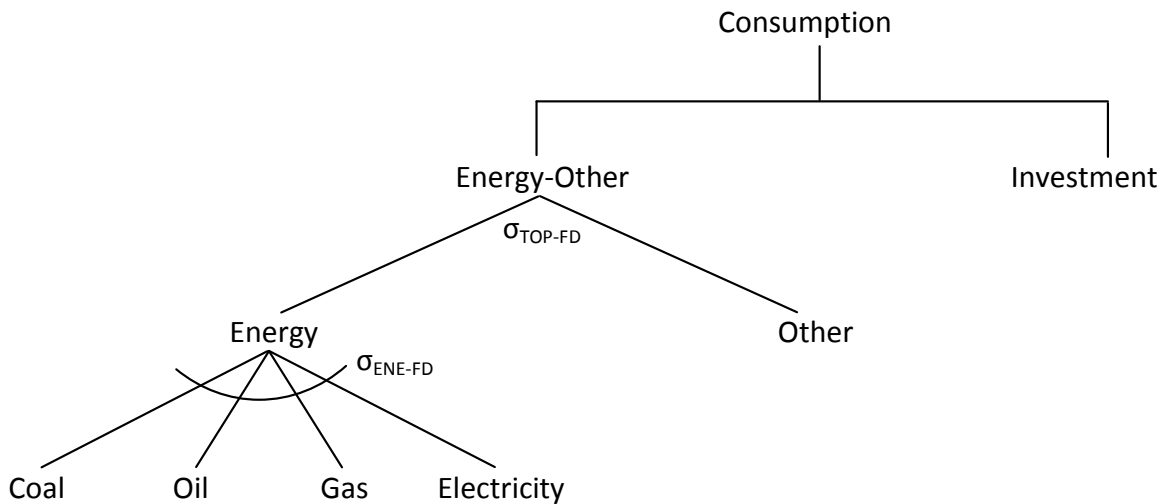
Figure 4.6 Consumer Utility Function

4.1.3 Model Dynamics

The CGE model is dynamic, running from 2010 to 2030 in 5-year time steps. 2010 is the benchmark year that is calibrated to the SAM data. The model then solves for 2015-2030. The processes that govern the evolution of the economy and its energy characteristics over time are: (1) capital accumulation, (2) fossil fuel resource depletion, (3) availability of low-carbon electricity technology, (4) population growth, and (5) energy efficiency improvements. The first three processes are endogenous while the last two are exogenous.

(1) Investment and capital accumulation:

An investment sector is specified that produces an aggregate investment good equal to the level of investment determined by the consumer utility function. The accumulation of capital is calculated as investment net of depreciation. The capital depreciation rate used is 7% per year.

Of particular importance for the uncertainty work is capital vintaging, which is applied to the electricity sector and reflects the irreversibility of decisions. Once a power plant is built, it is expected to operate as part of the electricity system for a long time, typically over 40 years. Therefore it can take a long time to change the composition of the electricity sector. This irreversibility is represented in the model by capital vintaging. Capital vintaging tracks the amount of electricity generation capacity available from previous years, remembering for each "vintage" (i.e. time period of installation) the technical features of the capacity (e.g. amount capital vs. labor vs. fuel, etc.). In the electricity sector, the model distinguishes between malleable and non-malleable (vintage) capital. New capital installed at the beginning of each period starts out in a malleable form. At the end of the period a fraction of this capital (70%) becomes non-malleable and frozen into the prevailing techniques of production. As the model steps forward in time it preserves four vintages of rigid capital (with 5-year time steps, this is 20 years), each retaining factor input shares at the levels that prevailed when it was installed with no possibility of substituting between inputs (i.e. elasticities of substitution equal to zero). In this way, today's decisions about how much of each technology to build affect the electricity system long into the future.

While vintage capital remains available until it depreciates away, if there is no demand for it then it may be unused. In the stochastic modeling approach employed in this dissertation, there are resolutions of the future that are different than what was expected when the capital was put in place. As a result, circumstances can arise in which even though the capital costs are sunk, the variable costs of vintage production are greater than the full cost of investing in new generation from an alternative technology, in which case vintage capital would go unused.

(2) Fossil Fuel Resource Depletion:

To capture major long-run dynamics of resource prices, a simple resource depletion model is included. An initial resource base (R) is defined and the endowment of fossil fuel

resources over time is subject to depletion based on physical production of the fuel in previous periods:

$$\text{endowment}_{ff,t} = \text{endowment}_{ff,0} * (R - \sum_{tt<t} \text{production}_{ff,tt}) / R \qquad \text{(EQ. 4.4)}$$

(3) Availability of Low-Carbon Electricity Technology:

As stated above, the low-carbon technology is not initially used and only enters if and when it becomes economically competitive. As noted by Jacoby *et al.* (2006), observations of penetration rates for new technologies typically show a gradual penetration, for which there are numerous reasons, including limited trained engineering and technical capacity to install/operate these technologies and electricity system adjustment costs. To reflect this trend, a fixed factor resource is included in the model which acts as an adjustment cost to the expansion of a new technology. The fixed factor component can be thought of as the inverse of a resource depletion process. Initially, a very small amount of fixed factor resource is available. Once there is installation of the new capacity, the fixed factor resource grows as a function of the technology's output in the previous period, on the basis that as production capacity expands the capacity to produce and install the technology also expands. The fixed factor is a necessary input as shown in Figure 4.5, and it limits low-carbon capacity expansion in any period by the amount of the fixed factor resource available in that period, subject to the ability to substitute other inputs for it. As low-carbon output expands over time, the fixed factor endowment is increased, and it eventually is not a limitation on capacity expansion. The intuition is that expansion of output in period t incurs adjustment costs, but the experience gained leads to more engineering and technical capacity in period t+1. In the model the fixed factor (FF) function is defined to grow quadratically (consistent with Reilly *et al.*, 2012; Paltsev *et al.*, 2005; McFarland *et al.*, 2004):

$$FF_{r,l,t+1} = FF_{r,l,t} + SH * (\alpha Q_{r,l,t} + \beta Q_{r,l,t}^2) \qquad \text{(EQ. 4.5)}$$

where:
Q = output,
l = technology,
SH = initial input share of fixed factor in production (5%),
α and β = coefficient parameters.

The α and β parameters were estimated by evaluating annual additions to nuclear capacity as a function of total output of nuclear electricity for the 1970's through the mid-1980's in the US

and France on the basis that was an analog situation to what we might expect with other advanced technologies in their early introduction (Reilly *et al.*, 2012). The resulting values are: α= 0.93 and β=3.2. The initial input share of fixed factor in production is assumed to by 5%.

(4) Population Growth:

In the model, population growth is assumed equal to growth in labor supply. Labor (L) grows according to a defined growth rate (gr) of 3% per year (the rate is raised to the power of 5 due to the five year time steps):

$$L_{t+1} = L_0 (1 + gr_t)^5 \qquad \text{(EQ. 4.6)}$$

(5) Autonomous Energy Efficiency Improvements:

Autonomous Energy Efficiency Improvements (AEEI) is a stylized way to represent the evolution of non-price induced, technologically-driven changes in energy demand. AEEI is an exogenous time-trend that reduces the amount of energy required (EI) in a sector (s) to produce the same amount of output. The efficiency improvement rate (aeei) is defined in the model to be 1% per year.

$$EI_{s,t+1} = EI_{s,0} / (1 + aeei)^5 \qquad \text{(EQ. 4.7)}$$

4.1.4 Policy

This dissertation considers policies that place a cap on the total amount of emissions from the electricity sector. When a cap is placed on emissions, the model calculates a shadow price for the constraint that can be interpreted as the price that would be obtained in a permit market under a cap-and-trade system. A cap-and-trade policy is modeled by defining a limited number of permits available, and requiring one permit for the emission of one unit of CO_2. The use of coal and natural gas in electricity must be accompanied by enough permits to cover the resulting emissions. Because permits are a limited resource, there is a market and a resulting price for them (the carbon price). This price represents the marginal costs of emissions reductions. In the model, the revenue from the permits is distributed back to households (the representative consumer). The policy cost is measured as the change in total consumption compared to when there is no policy.

4.1.4 Model Solution and Outputs

Equilibrium is solved for as a mixed complementarity problem (MCP) (Mathiesen, 1985; Rutherford, 1995).[8] Section 4.1 describes the optimality conditions that must be met in order to find equilibrium. The MCP formulation solves a system of equations for a set of prices and quantities such that all of the optimality conditions are met. In this way, the MCP approach finds a market solution rather than a true welfare maximum solution (which is only possible under idealized market conditions). However, this is an advantage as we are more interested in how the market economy responds than an idealized solution.

The results of the model outputs include consumption, policy cost, emissions, emissions price if there is a policy, commodity prices (oil, coal, natural gas, electricity, etc.), and the energy and electricity mix. This work focuses on the electricity mix, emissions, and consumption/policy costs. The results from the model depend on a number of aspects of model structure and particular input assumptions. For example, the difficulty of achieving any emissions policy is influenced by assumptions about population and energy efficiency improvements that underlie the no-policy reference case. The elasticities of substitution are critical assumptions about the relative ease of substitution among the inputs to production and the behavior of consumers in the face of changing prices.

4.2 DP-CGE Model

The CGE model described above is reformulated as a two-stage stochastic dynamic program (DP) to create the DP-CGE model. The deterministic CGE model is a myopic recursive-dynamic model that solves for each time period sequentially. For a given period, the original CGE model chooses an electricity technology mix (and all other outputs) based on the current-period maximization of consumption. However, the objective of this research is to find the technology mix in each period that maximizes the current period consumption *plus* the *expected future* consumption. Dynamic programming is utilized to do so.

---

[8] Solving for equilibrium as a constrained optimization problem works for simple problems, but the usefulness of that approach breaks down quickly as the model becomes bigger and more complex in its representations. The algebra becomes tedious and challenging when solving for multiple markets and constraints at once.

4.2.1 DP Formulation

There are several key components of the DP formulation: the temporal sequencing, the decision-maker, the objective function, the decisions, the uncertainties, and the states. This section describes each of those components.

(1) Temporal Sequencing:

The DP is framed as a two-stage finite horizon problem. The underlying CGE model continues to run in 5-year time steps, but the time horizon is divided into two decision stages for the DP. Stage 1 includes CGE periods 2015 and 2020 while Stage 2 includes 2025 and 2030 (and 2010 is the benchmark year). In each stage, the DP decisions are made for the two CGE periods included in that stage.

(2) Decision-Maker:

In the underlying CGE model, the decision-maker is a hypothetical central planner of the economy. Although the optimal electricity mix is solved as if from the perspective of a central planner, one can think of it as the aggregate result of individual and identical firms maximizing their own profits according to their production functions, input costs and the policy constraints imposed by the central planner.

(3) Objective Function:

The objective is to choose actions to maximize total expected discounted social welfare in the economy over the planning horizon. In terms of the Stage 1 (near-term) decisions, the goal is to maximize current period consumption plus discounted expected future consumption. Utilizing the Bellman equation, defined in Chapter 3, the objective function is:

$$V_t = max_{x_t}[C_t(S_t, x_t) + \gamma E\{V_{t+1}(S_{t+1}(S_t, x_t, \theta_t))\}] \qquad \text{(EQ. 4.8)}$$

where:
t is decision stage,
V is total value,
S is state (electric power capacity level of each technology and cumulative emissions level),
C is economy-wide consumption (welfare),
x is decision set (low-carbon share of new electricity and amount of emissions reductions),
θ is uncertainty (probabilities assigned to Stage 2 policy and low-carbon technology cost)
γ is discount factor = (1-discount rate). Discount rate = 4%.

(4) Decisions:

In the DP, two decisions are made (so x_t in Equation 4.8 is a vector with two elements). The first decision is the low-carbon technology's share of *new* electricity in each stage. The decision about the share of the low-carbon technology generation only applies to new electricity generation. This reflects the real-world question: in expanding electricity capacity, how much of the new capacity should consist of low-carbon technologies? This decision is exogenously imposed on the CGE model.

The second decision is Stage 1 reductions of electricity emissions. In anticipation of policy in Stage 2, it may be desirable to begin reducing emissions in Stage 1 to help reduce the cost of meeting the expected future policy. This decision to reduce Stage 1 emissions via a "self-imposed" emissions cap provides a price signal in the CGE model that affects the operation of existing electricity capacity as well as the optimal share of the low-carbon technology in new electricity. If the share of the low-carbon technology in new electricity was the only decision, investing in new low-carbon generation would be the only way to reduce electricity emissions in Stage 1. However, near-term emissions can also be reduced by changing the way existing capacity is operated, particularly by shifting away from coal and toward natural gas generation and/or leaving some vintage capacity to go unused or underutilized. Including this second decision in the DP provides a price signal (the shadow price (i.e. carbon price) of the self-imposed emissions constraint) for the CGE model to endogenously react to by changing the operation of vintage capacity to reduce near-term emissions. Ultimately this emissions reduction decision variable affects choices of coal vs. natural gas, conventional vs. low-carbon, and building new vs. operating existing capacity differently.

(5) Uncertainties:

In the DP two uncertainties are investigated. First, uncertainty in the Stage 2 emissions cap policy is explored, and then uncertainty in the Stage 2 cost of the low-carbon technology. The potential policies are defined as caps on the cumulative emissions from the electric power sector from 2015 to 2020. The cost of the low-carbon technology is defined relative to the cost of conventional generation. For each of these uncertainties three possible realizations are defined and assigned probabilities. Details about the uncertainties and their implementation are discussed in Section 4.3.

(5) State Variables:

In the DP framework state variables are used to make the system Markovian—all information needed to make a decision is contained in the current state so that no further information is needed. There are two key state variables for this problem. First is the installed capacity level for each generation technology, which informs the decisions about operating existing capacity and building new capacity. Second is the cumulative emissions level, which also affects expansion and operation decisions in order to meet future policy. In the implementation, additional state information must be tracked within the CGE model, such as the level of capital stock, labor, natural resources, fixed factor resources, and energy conversion efficiency.

The DP-CGE model is schematically illustrated with the representative decision tree diagram in Figure 4.7, using uncertainty in the Stage 2 emissions cap policy as an example. The DP-CGE model is solved in two steps. First, the CGE model is run for each stage for each possible scenario (each combination of decision and uncertainty realization), calculating the total consumption for each stage. The decisions and uncertainty realizations are exogenously imposed on the CGE model, which then endogenously chooses all other output quantities, including the shares of natural gas and coal generation. Second, backward induction is performed by the DP in Matlab using the consumption values for each stage and the probabilities (p1, p2, p3) of the uncertainty realizations. The DP assumes that the true policy is learned before Stage 2 decisions are made, so for each possible emissions cap the optimal Stage 2 decision is determined. The consumption resulting from each of the three optimal Stage 2 decisions (one for each possible emissions cap) is then multiplied by the probability of each cap and summed together to obtain an expected value of Stage 2 consumption. Then, "folding back", a single optimal Stage 1 decision is identified that maximizes Stage 1 plus expected Stage 2 consumption. In effect, the CGE model performs intra-period optimization and the DP performs inter-period optimization.

Figure 4.7 Decision Tree Depiction of DP-CGE Framework

4.3 Characterization of Uncertainties

4.3.1 Uncertainty in Future Policy

The situation with regard to policies that may limit GHG emissions and/or provide incentives for alternative electricity technologies in the U.S. is particularly uncertain now due to the highly uncertain economic environment, the tumultuous political environment, and a huge deficit in federal spending. Even though the science of global climate change is firmer than ever in suggesting significant risks and implying a need to transform the global energy system, the political support in the U.S. for undertaking federal policy to push the U.S. system in that direction is probably weaker than it has been since 1997 when the Senate voted unanimously against the Kyoto concept of binding targets.

In the short-run, the prospects for policy are low. There is political gridlock in Congress on the issue and opinion polls of the American public show that while a majority still believes in the science of climate change, the appetite for action is weak. The authors of a poll conducted by

the Yale University Project on Climate Change Communication (Leiserowitz *et al.*, 2010) conclude that widespread misconceptions "lead some people to doubt that climate change is happening or that human activities are a major contributor, to misunderstand the causes and therefore the solutions, and to be unaware of the risks." Thus a majority of the public agreeing that climate change is real does not translate to broad support for aggressive action because beneath the weak grasp of the issue is a lack of knowledge or unwillingness to understand the source or seriousness of the risk. Far higher on Americans' priority list are the very high unemployment rates and concern about the federal deficit. Those issues have dominated recent elections and another poll asking about Americans' concerns and priorities had the climate issue last among 25 issues identified. As a result, over the next 5 years at least, R&D funding, tax incentives, and other direct spending on energy (or any other budget priority) are on a collision course with the much bigger forces of the out of control deficit, the battle over tax cuts, and the recession. However, there is still the looming possibility that the EPA will impose a GHG cap-and-trade system or other emission regulations itself.

In the longer-run (10-50 years), there is much greater uncertainty about whether policy will be implemented and its potential stringency. This uncertainty is driven by two main forces: political uncertainty and scientific uncertainty. On the political front, it is unclear whether government leaders will be able to come to a consensus on this issue and implement a policy. If the political will can be garnered, it seems likely that policy would take a form similar to the Waxman-Markey cap-and-trade bill since that approach had fairly broad industry support including support from utilities.

Scientific uncertainty also plays an important role. While the science that climate change poses a significant risk is well-established, the science linking emissions reductions to climate outcomes is far less certain. At the global level, long-term climate goals are discussed in terms of stabilizing atmospheric concentrations of CO_2 or GHG emissions at certain levels. Different targets are discussed, such as stabilization at 450 parts per million (ppm) or 550 ppm. These long-term targets require differing levels of emission reduction policy. The problem is that it is unclear which target to aim for, and therefore the level of emissions reductions to pursue, because the science is uncertain about how different emission policy paths will impact the climate system. An extensive study by Webster et al (2012b), demonstrates this uncertainty. Figures 4.8 and 4.9, from that study, show the uncertain impact of five different emission policy

83

paths on CO_2 concentrations and global mean surface temperature change. The uncertainty ranges are significant. Thus while the science indicates the importance of taking some actions, it is difficult to determine what level of action to take. It is unknown how much of the scientific uncertainty will be resolved in coming years, and it is also unclear to what extent policy will respond to scientific information.

(Source: Webster et al., 2012b)
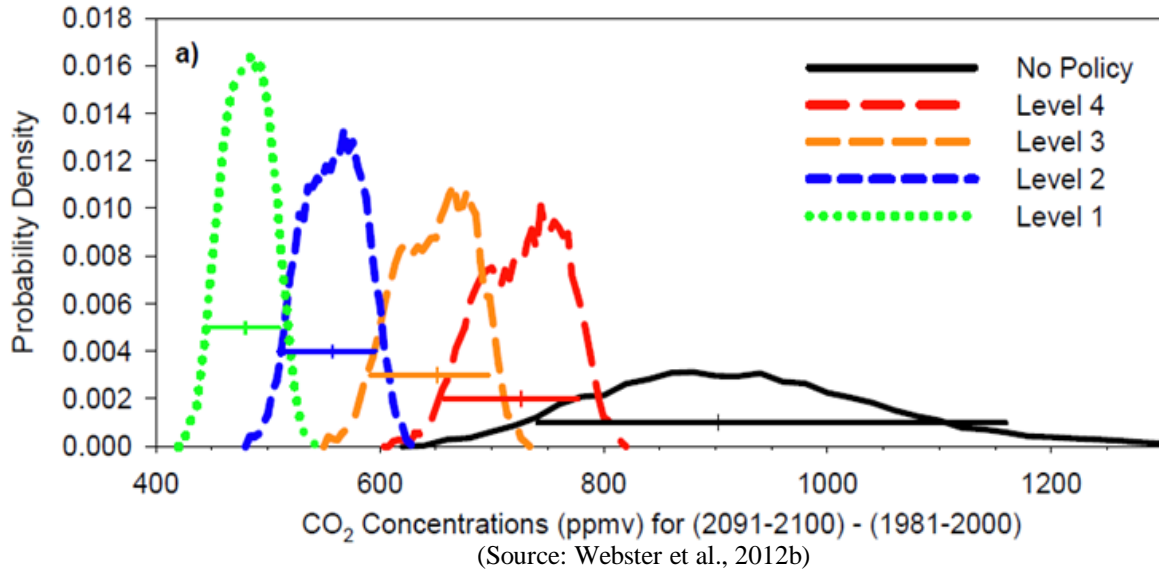
Figure 4.8 Impact of Emission Reduction Policies on CO_2 Concentrations
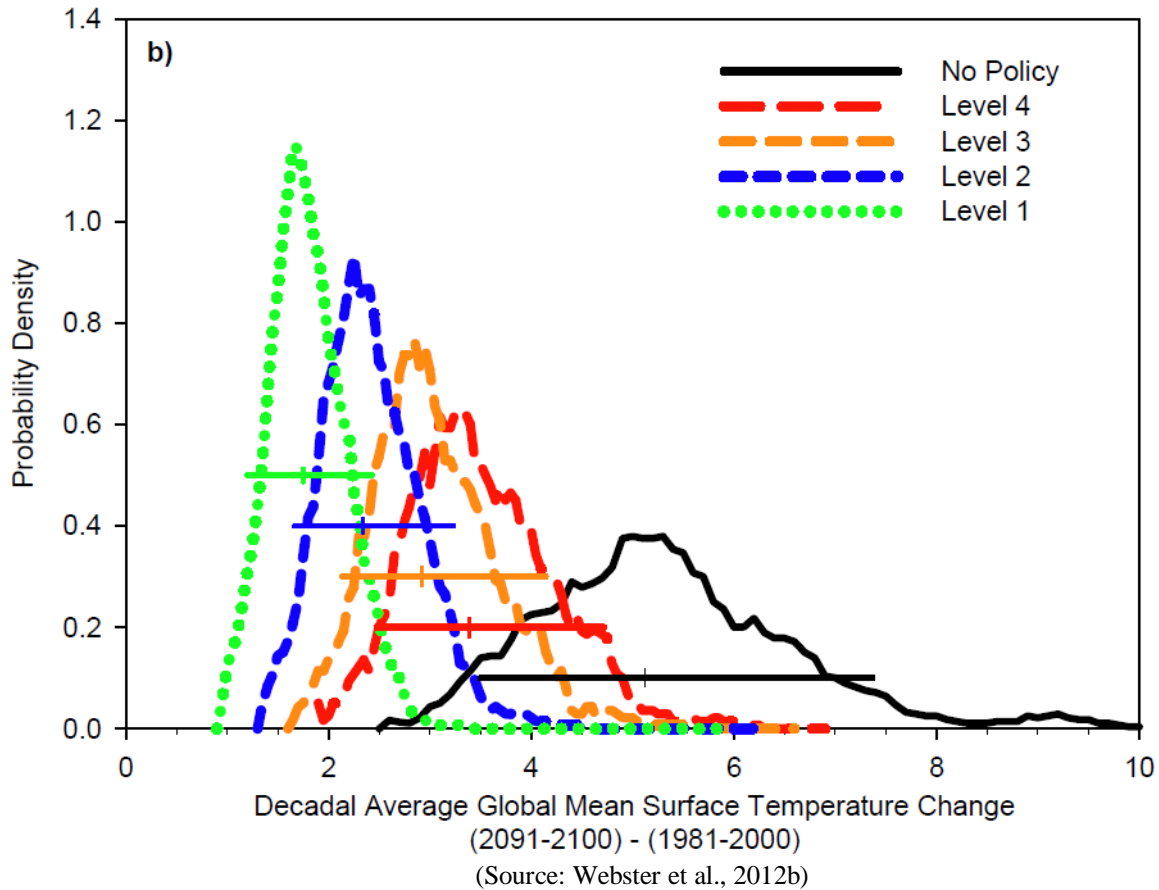
Figure 4.9 Impact of Emission Reduction Policies on Global Mean Surface Temperature

4.3.2 Representing Policy Uncertainty in the DP-CGE Model

The DP-CGE model uses a discrete approximation of the continuous uncertainty in future emissions policy. Specifically, a discrete three-point probability distribution with three policy scenarios is assumed: (1) no policy, (2) an emissions cap of 20% below no policy emissions (-20% Cap), and (3) an emissions cap of 40% below no policy emissions (-40% Cap). Each of these scenarios is assigned an associated probability, which collectively sum to one. In this analysis, the emissions caps apply to *cumulative* emissions from 2015 to 2030 from the *electricity sector*. Reference no policy cumulative emissions are determined by running a no policy case and summing the resulting electricity emissions over the period 2015-2030. The other two scenarios then require cumulative electricity emissions to be either 20% or 40% below these cumulative reference emissions. The policy cases are focused on cumulative emissions because it is cumulative emissions, not the specific emissions path over time, which matter most

for climate change. This framing is also consistent with past cap-and-trade proposals in the U.S., which included intertemporal flexibility (banking and borrowing of emission permits) in meeting the emissions target.

To provide some context for these caps, according to analysis of the Waxman-Markey cap-and-trade bill, by 2030 the bill would either result in 27% or 41% cumulative reductions of electricity emissions, depending on the assumption about the availability offsets to help meet the cap (Paltsev *et al.*, 2009). Another point of comparison is President Obama's recently proposed climate plan, which includes reducing economy-wide GHG emissions by 17% below 2005 levels by 2020. Taking that goal and linearly extrapolating to 2030 allows calculation of cumulative emissions from 2015 to 2030. Comparing those cumulative emissions to the reference cumulative emissions from this model results in cumulative emissions reductions of about 33%. Given that those are reductions in economy-wide emissions, and one would expect the electricity sector to reduce more than other sectors (due to cheaper abatement options), it is reasonable to suspect that under the Obama plan cumulative reductions in electricity sector emissions would be similar to the 40% cap case.

In the DP-CGE model, the policy is assumed to begin in 2025, the start of Stage 2. Following the classic act-then-learn framework, Stage 1 decisions about electricity technologies and emission reductions must be made without knowing which of the three policies will be in effect in 2025, but with expectations about which policies are most likely. The expectations are reflected in the probabilities assigned to each of the possible policies. Before Stage 2 decisions are made, the policy is revealed, so Stage 2 decisions are made with perfect information about the emissions limit.

The perceived probabilities of policies are difficult to ascertain. This work takes two approaches to assigning expectations to the policies. First, it explores a wide range of possible discrete three-point probability distributions and investigates the effect the differing expectations have on the optimal decisions. Second, a small survey is conducted to get a sense of the expectations about future policy that real world investors might use to make their decisions. These approaches of assigning probabilities are laid out in Chapter 5.

4.3.3 Uncertainty in Future Technology Costs

The future costs of technologies are highly uncertain, particularly for advanced low-carbon technologies that have not yet penetrated the market but may in future years. As with estimates of current costs (e.g. LCOE, capital and O&M costs), there is a wide range of estimates of future technology costs. However, a range of point estimates is a poor proxy for an uncertainty range and tends to be much narrower. An alternative way to estimate uncertainty is through expert elicitation of probability distributions. Although expert judgments that can be biased (Tversky & Kahneman, 1974; Morgan & Henrion, 1990), if protocols (e.g. Staël von Holstein & Matheson, 1979; Morgan & Henrion 1990) are followed to minimize bias, this approach can be an effective way of defining uncertainty ranges for future technology costs.

Several recent studies have conducted expert elicitations of the costs of advanced generation technologies. Anadoni *et al.* (2012) obtained probabilistic estimates of the 2030 cost and performance of nuclear generation technologies under different scenarios of government RD&D spending from 30 U.S. and 30 European nuclear technology experts. As an example of their results, Figure 4.10 shows each expert's best estimates (black circles) and 90th and 10th percentile error bands (vertical lines) for Generation IV nuclear capital costs in 2030 under a business as usual RD&D funding scenario. Bosetti *et al.* (2012) conducted a similar study eliciting 16 European experts for their probabilistic estimates of the 2030 costs of solar technologies under different RD&D spending scenarios (Figure 4.11). Both the nuclear and solar figures show that each expert perceives a different uncertainty range. Similar elicitations have been or could be conducted for other advanced technologies, such as CCS (e.g. Baker *et al.*, 2009; NRC, 2007; Chan *et al.*, 2010) and wind. There are also approaches for compiling individual expert predictions into a single probability density function (PDF) to express the range of uncertainty (Genest & Zidek, 1986; Clemen & Winkler, 1999). For example, Chan *et al.* (2010) use a Monte Carlo analysis weighting each expert's cost distribution equally to generate a single cumulative distribution function of the 2030 capital costs for coal with CCS for four different scenarios (Figure 4.12).
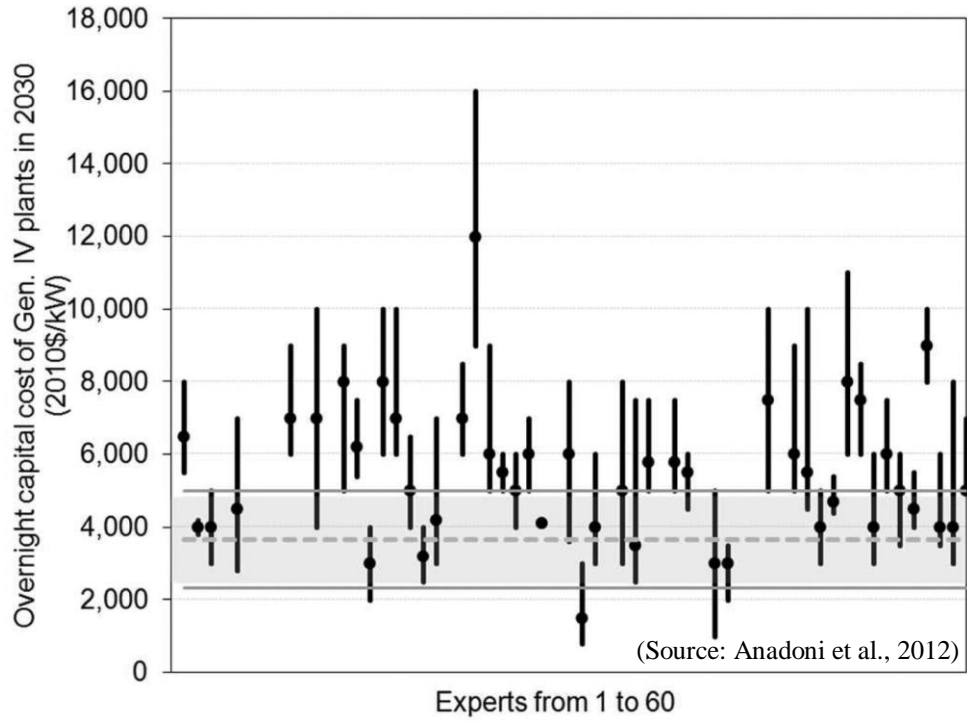
Figure 4.10 Expert Elicitation of Nuclear Capital Costs in 2030
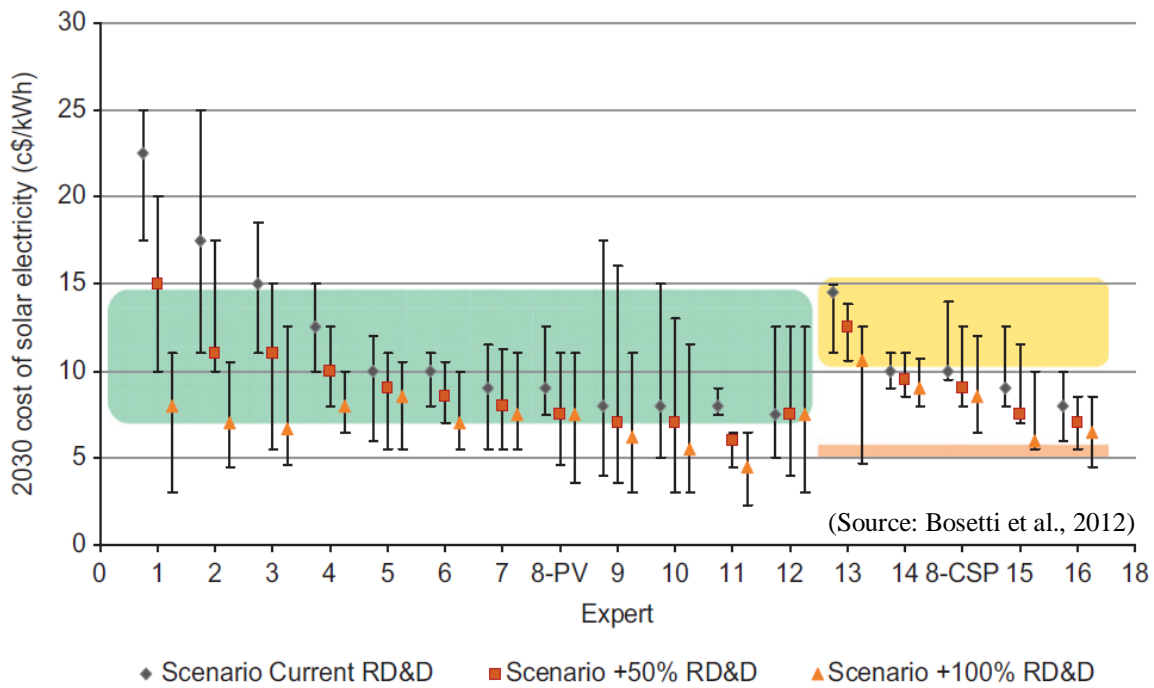
Figure 4.11 Expert Elicitation of Solar Electricity Cost in 2030
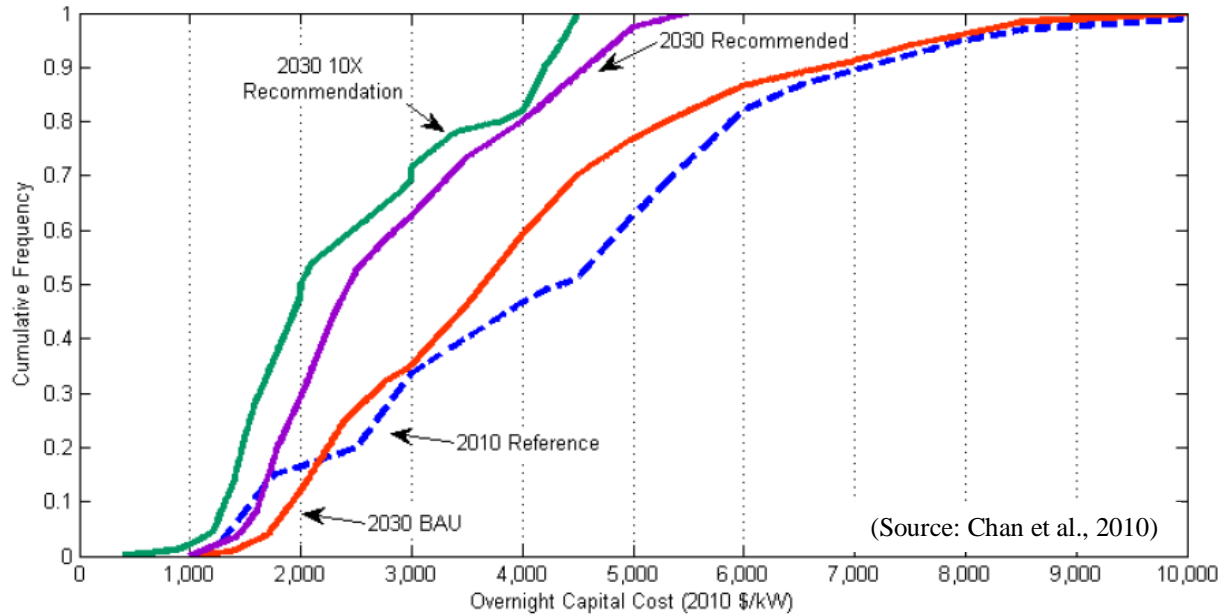
Figure 4.12 Cumulative Distribution from Combined Expert Elicitation of Coal with CCS Capital Costs in 2030

These studies demonstrate the wide uncertainty range of the future costs of generation technologies. In the face of this uncertainty, decisions must be made about which electricity technologies to build. When looking at a particular technology one must not only look at its cost, but also its cost relative to the other technologies available.

4.3.4 Representing Technology Cost Uncertainty in the DP-CGE Model

In the DP-CGE model, Stage 1 decisions about electricity technologies and emission reductions must be made without knowing which of the three low-carbon technology costs will be realized in Stage 2, but with expectations about which costs are most likely and an understanding that the cost will be driven by near-term investments in low-carbon technology. Before Stage 2 decisions are made, the low-carbon technology cost is learned.

Within the CGE model, the decision of which technologies to build is driven by the relative costs of the technologies. Recall from Chapter 4 that the cost of the low-carbon technology is initially set in the model by a markup, which is the cost relative to the conventional generation against which it competes. All inputs are multiplied by the markup. The markup is initially set to 1.5, indicating that the low-carbon technology is 50% more expensive than conventional electricity in the base year of the model. To model uncertainty in the Stage 2 cost of

the low-carbon technology, the markup is also set at the beginning of Stage 2 and its value made uncertain. Because the low-carbon technology represents multiple advanced low-carbon technologies (e.g. wind, solar, nuclear, CCS, etc.), the Stage 2 markup represents the markup of the cheapest of the low-carbon technologies, whichever that ends up being.

The DP-CGE model uses a discrete approximation of the continuous uncertainty in future cost of the low-carbon technology. Specifically, a discrete three-point distribution is assumed with three low-carbon technology cost scenarios: (1) a markup of 1: low-carbon generation costs the same as conventional generation (MU1), (2) a markup of 1.5: the low-carbon generation continues to cost 50% more than conventional generation (MU1.5), and (2) a markup of 3: the low-carbon generation costs triple conventional generation (MU3). These markup values are informed by the expert elicitation studies discussed above. Each of those studies elicits technology costs for 2030 (which can be used for the Stage 2 markup in the DP-CGE model). Looking at the cumulative probability distribution for CCS capital costs from Chan et. al (2010) (Figure 4.12), we can estimate the 5^{th}, 50^{th} and 95^{th} percentiles. Putting the capital costs numbers into an LCOE calculation (Morris et. al 2010) provides markup values for those percentiles. The 5^{th}, 50^{th} and 95^{th} percentile markups are 1.06, 1.63, and 2.82. Baker et al (2009) also provides a combined probability distribution of the expert elicitations for CCS. Markups estimated from figures in that paper are 1.075, 1.35, and 1.45 for the 5^{th}, 50^{th} and 95^{th} percentile. For other studies that do not provide a cumulative probability distribution, we can look at the probability range of expert judgments and translate into markups using the Morris et. al (2010) LCOE calculation. For nuclear, estimating values from Figure 4.10 (and removing the high outlier) gives markups ranging from 0.68 to 4.01. For solar, estimating values from Figure 4.11 gives markups of 0.67 to 4.18.

Table 4.3 Markups Derived from Expert Elicitation Studies

	Chan *et al.* (2010)	Baker *et al.* (2009)	Bosetti *et al.* (2012)	Anadon *et al.* (2012)
	CCS	**CCS**	**Solar**	**Nuclear**
5th percentile OR Min	1.06	1.075	0.67	0.68
50th percentile	1.63	1.35		
95th percentile OR Max	2.82	1.45	4.18	4.01

Note: The percentiles apply to both CCS estimates while the Min and Max apply to solar and nuclear.

With these studies in mind (summarized in Table 4.3), it is assumed that markups of 1, 1.5 and 3 are reasonable approximations for the 5[th], 50[th] and 95[th] percentiles for the cost of the general low-carbon technology for this model. Using the extended Pearson-Tukey discrete approximation method (Keefer & Bodily, 1983), the base probabilities of high (MU3), medium (MU1.5) and low (MU1) outcomes are assigned to be 0.185, 0.63, and 0.185 respectively. Additional mean-preserving probability spreads are also explored. Stochastic technological learning is then incorporated by having the Stage 1 low-carbon generation shares determine the probabilities of the Stage 2 low-carbon technology cost scenarios. The Pearson-Tukey probabilities are used if there is no low-carbon investment in Stage 1. According to a learning rate, as the amount of Stage 1 low-carbon generation increases, the probability of low Stage 2 markup increases and the probability of a high Stage 2 markup decreases. Technological learning rates and sensitivity to the rates are presented in Chapter 7.

4.4 DP-CGE Implementation

To numerically implement the DP-CGE model two important changes were made to the CGE structure and the decision and uncertainty spaces were discretized. This section discusses the implementation.

In terms of structural changes to the CGE model, first it is transformed into a two-stage model that starts and stops at each stage and solves using saved state variables instead of continuously looping through each time period. Stage 1 is run and at the end the current levels of all state variables are saved. Then at the start of Stage 2, the Stage 1 levels of all the variables are read in as the Stage 2 base year levels and Stage 2 is run from there. Within each stage, the model loops over 5-year time periods in the conventional sequential manner. This setup allows Stage 2 to be run independent of Stage 1. Thinking of a decision tree, we can run a number of Stage 2 scenarios that branch from a single Stage 1 scenario without having to rerun Stage 1 first each time. This setup also allows Stage 2, as well as Stage 1, scenarios to be run in parallel.

The second structural change is to set the low-carbon technology share exogenously, removing the technology mix decision from the CGE model. The CGE model would choose an electricity mix based only on current-period consumption and would not consider the expected future consumption. To get around this, different technology mixes are forced, requiring the model to explore a wide range of mixes that can then be evaluated based on their contribution to

the two-stage Bellman optimization. Specifically, the share of new electricity from the low-carbon technology in both Stage 1 and Stage 2 is exogenously forced. Given the share of low-carbon generation, the CGE model then determines the optimal share of coal vs. natural gas generation as it solves the equilibrium for the economy. To implement this low-carbon technology share forcing, the production functions for conventional electricity and low-carbon electricity are altered to produce two distinct outputs ("conventional electricity" and "low-carbon electricity") instead of both producing the identical output ("electricity"). Then a "transformation" sector is added to the model that produces "electricity" by using "conventional electricity" and "low-carbon electricity" as inputs in fixed proportions which depend on the share of low-carbon technology being forced (Figure 4.13).

Electricity

Backstop Electricity   Fossil Electricity

Figure 4.13 Electricity Production with Forced Low-Carbon Technology Share

A stochastic dynamic program, which includes decision-making under uncertainty with learning, is then wrapped around the restructured two-stage CGE model.

Within the DP, the continuous decision space is discretized for numerical implementation. For the decision of low-carbon generation's share of new electricity, the DP explores a range of 0% to 50% low-carbon generation in steps of 5% in Stage 1, and a range of 0% to 80% in 5% steps in Stage 2.[9] The Stage 1 low-carbon technology share applies to 2020 and the Stage 2 share applies to 2030—for example, there must be 15% low-carbon generation by 2020 and 40% low-carbon generation by 2030. For 2025 the low-carbon technology share imposed is the average of 2020 and 2030 shares. For 2015 the low-carbon technology share

---

[9] These ranges were chosen after extensive testing of the model under a wide range of scenarios found that Stage 1 low-carbon share never rose above 50% and Stage 2 never rose above 80%.

imposed is the minimum of 5% and half of the 2020 share. This was determined after extensive testing of the model showed that 5% low-carbon generation overcomes the fixed factor resource constraint (i.e. enough low-carbon generation has been built that the fixed factor resource grows to the point that it does not limit additional low-carbon investment). Because all low-carbon generation built in 2015 is constrained by the fixed factor (since 2015 is the first year the low-carbon technology can enter), it is never economically optimal to build more than 5% in 2015. If you want to build a lot of low-carbon generation by 2020, the best strategy is to only build 5% in 2015 to overcome the fixed factor constraint and then build the desired amount in 2020 when the fixed factor constraint is no longer binding.

The second decision about Stage 1 reductions of electricity emissions is also discretized. The DP explores a range of Stage 1 emission growth rates (from 0.8 to 1.2 in 0.05 steps) that apply to each time period within Stage 1. Applying the growth rate to 2010 base year emissions creates an emissions cap that is imposed in 2015, and applying the growth rate to the 2015 emission cap creates a cap that is imposed in 2020. Those caps are then implemented in the CGE model. These are the "self-imposed" emissions caps described in the previous section. Table 4.4 shows what each growth rate means in terms of cumulative Stage 1 emissions reductions relative to the base case in which no reductions are made. A growth rate of 1.2 is equivalent to making no Stage 1 emissions reductions.

Table 4.4 Stage 1 Emissions Reduction Decision Space

Emissions Growth Rate	Percentage Reduction in Cumulative Stage 1 Emissions Relative to Reference
0.800	45%
0.850	40%
0.900	35%
0.950	30%
1.000	24%
1.050	18%
1.100	13%
1.150	6%
1.200	0%

The uncertainty space is also discretized. For both the Stage 2 policy and low-carbon technology cost three possible realizations are defined and assigned probabilities. The cumulative cap policies are enforced by taking Stage 1 cumulative emissions and subtracting that from total allowed cumulative emissions, which gives the amount of cumulative emissions allowed in Stage 2. That amount is then divided by two and applied as the cap in 2025 and 2030. A no policy case is represented by a cumulative emissions cap equal to reference cumulative emissions, and therefore nonbinding. The low-carbon technology cost is implemented using a multiplicative markup factor which describes the cost of low-carbon generation relative to conventional generation.

The overall size of the DP-CGE model is determined by the number of decision periods, the decision space and uncertainty space. Taking uncertain policy as an example, at the level of discretization described, there are a total of 6,171 possible Stage 1+Stage 2 path scenarios. The CGE model is run to generate results for each of the scenarios and the DP uses that information together with the probabilities of the cap policies to find the optimal Stage 1 decision set.

The main goal of this simple DP-CGE model is to provide a numerical modeling framework for decision-support under uncertainty, and to understand if and how the electricity strategy under uncertainty differs from a deterministically designed strategy. This is explored in the chapters that follow.

Chapter 5: Model Results: Policy Uncertainty

This chapter uses the DP-CGE model described in Chapter 4 to investigate near-term electricity investment and emissions reduction decisions under uncertainty in future climate policy. Section 5.1 gives an overview of the analysis that follows. Section 5.2 presents results from the model when the policy is known for certain. The results when the policy is uncertain are then presented in Section 5.3. A small survey and discussion of real world expectations about policy is presented in Section 5.4. Analyses of the cost of uncertainty and the value of including uncertainty follow in Section 5.5. Section 5.6 presents results from a sensitivity analysis investigating the impact of assumptions about limits to low-carbon generation growth rates. Finally, Section 5.7 provides a concluding discussion, and motivates the analyses in the remaining chapters.

5.1 Introduction

As discussed in Chapter 4, the DP-CGE model uses a discrete three-point probability distribution with three policy scenarios: (1) no policy, (2) an emissions cap of 20% below no policy emissions (-20% cap), and (3) an emissions cap of 40% below no policy emissions (-40% cap). Each of these scenarios has an associated probability, which collectively sum to one. The emissions caps apply to *cumulative* emissions from 2015 to 2030 from the *electricity sector*. Reference no policy cumulative electricity emissions are determined by running a no policy case and summing the resulting electricity emissions over the period 2015-2030.[10] The other two scenarios then require cumulative electricity emissions to be either 20% or 40% below these cumulative reference emissions.

In the DP-CGE model, the policy is assumed to begin in 2025, the start of Stage 2. Following the classic act-then-learn framework, Stage 1 decisions must be made without knowing which of the three policies will be in effect in 2025, but with expectations about which policies are most likely. The expectations are reflected in the probabilities assigned to each of the possible policies. Before Stage 2 decisions are made, the policy is revealed, so Stage 2 decisions are made with perfect information about the emissions limit.

---

[10] Reference cumulative electricity emissions from 2015-2030 computed from the model are ~30% higher than the latest EIA estimate (EIA, 2013b).

The Stage 1 decisions are: (1) the share of new electricity investment from each technology, and (2) emissions reductions (relative to reference emissions). New electricity investments include truly new investments as well as retrofit investments (i.e. investments to change existing electricity capital).[11] In all of the scenarios investigated, new electricity investment turns out to be responsible for approximately 40% of total electricity generation by the end of Stage 1. This share of total electricity is mainly driven by the increasing demand for electricity, with different decisions about near-term investments only changing the share slightly. Accordingly, as an example, if the optimal Stage 1 decision is to have 20% of new electricity investment be in low-carbon technologies, then 8% (20% of 40%) of total electricity will come from *new* low-carbon generation capacity.

5.2 Results under Policy Certainty

First, it is useful to examine the results from this model under perfect information about the future emissions policy. That is, knowing which policy will be in place in Stage 2, what should we do in Stage 1? This representation is akin to having perfect foresight. Figure 5.1 shows the Stage 1 and Stage 2 decisions about the share of new electricity investments from each technology under three deterministic scenarios: (1) it is known for certain there will be no policy, (2) it is known for certain there will be a 20% cap, and (3) it is known for certain there will be a 40% cap.

If there will be no policy, there is no reason to invest in low-carbon technologies. For this scenario, the optimal fuel mix from conventional electricity generated from new investment in Stage 1 is 63% natural gas and 37% coal. Because there is perfect information, we also solve for a single optimal Stage 2 electricity mix, which is 81% natural gas, and 19% coal. The rising share of natural gas is the result of coal prices rising relative to natural gas prices. If there will be a 20% cap, it is optimal to start investing in low-carbon technologies in Stage 1 in anticipation of the Stage 2 policy. The optimal Stage 1 new electricity mix is 5% low-carbon, 65% natural gas and 30% coal. Given the parameterization of the fixed factor resource constraint (Equation 4.5 in Chapter 4) and the resolution of the decision space, a 5% low-carbon investment share is the

---

[11] As Chapter 4 describes, 70% of capital investments are vintaged (locked into place), but 30% remain malleable, providing the flexibility to reinvest that capital. Reinvestments that go back into the electricity sector can be considered retrofitting existing electric generation capacity.

minimum amount required to overcome this constraint on the rate of technology expansion, thereby providing flexibility to expand low-carbon penetration to the desired amount in Stage 2. The Stage 2 optimal new investment mix is 40% low-carbon, 50% natural gas and 10% coal. If the policy will be a 40% cap, it is optimal to start investing aggressively in low-carbon technologies in Stage 1 in order to lower the costs of meeting the policy in Stage 2 policy. The optimal Stage 1 new electricity mix is 35% low-carbon, 51% natural gas and 14% coal, and the optimal Stage 2 mix is 60% low-carbon, 30% natural gas and 10% coal (Figure 5.1). These results show the effect of information about the future policy on near-term decisions. When we know what the future emissions policy will be, we know the optimal mix of near-tem investment. Given the formulation of the model, with an emission limit of at least 20% some near-term investment in low-carbon technologies is optimal.[12]
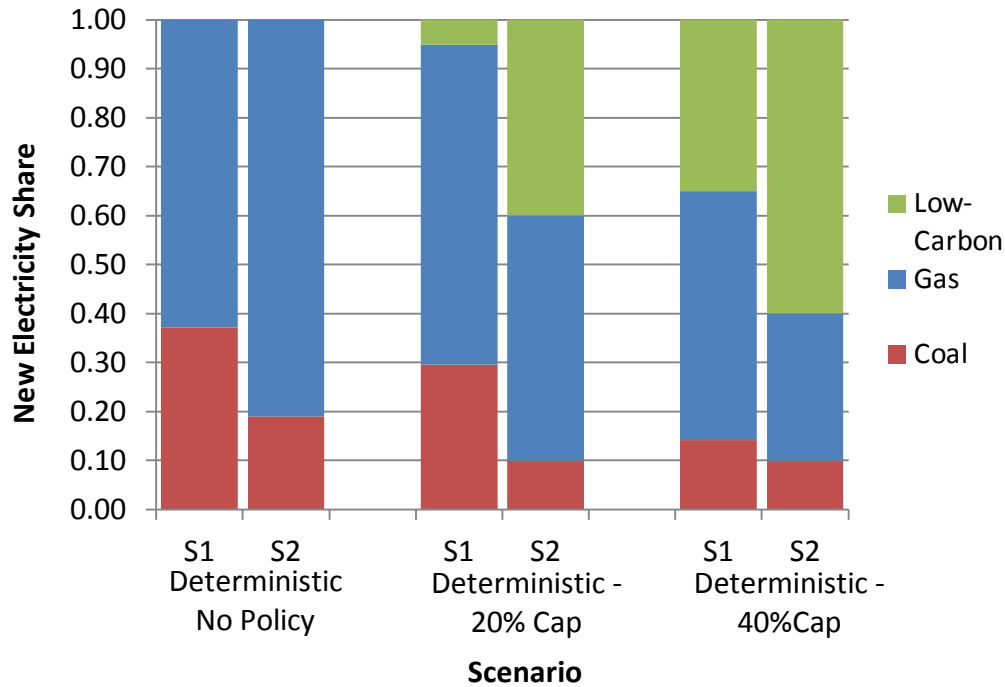
Figure 5.1 Stage 1 and Stage 2 Shares of New Electricity under Policy Certainty

---

[12] Emission limits below 20% may also warrant near-term investments in low-carbon technologies, but those scenarios were not explored in this work.

In Stage 1, decisions are also made about the level of emissions reductions. It is important to note that electricity emissions are the net effect of both new electricity investment decisions (i.e. which technologies to add to the mix) and decisions about how to operate existing (vintage) electricity capacity. Emission reductions can take place by adding new low-carbon capacity, by adding new non-coal conventional capacity, or by changing the operation of vintage capacity (for example to use less coal in existing capacity or less generation overall). Figure 5.2 shows the Stage 1 decisions about emissions reductions under the three deterministic scenarios. Emissions are shown relative to the no policy Stage 1 emissions. If we know there will be no cumulative emissions limit in Stage 2, emissions are equivalent to the reference case. If we know there will be a 20% cap, it is optimal to start reducing emissions by 6% below the no policy level in Stage 1 in anticipation of the Stage 2 policy. If we know there will be a 40% cap, it is optimal to reduce Stage 1 emissions by 24% below the no policy emissions. It is worth it to bear the costs of the near-term reductions in order to avoid having to make as drastic and costly cuts in Stage 2 to meet the policy.
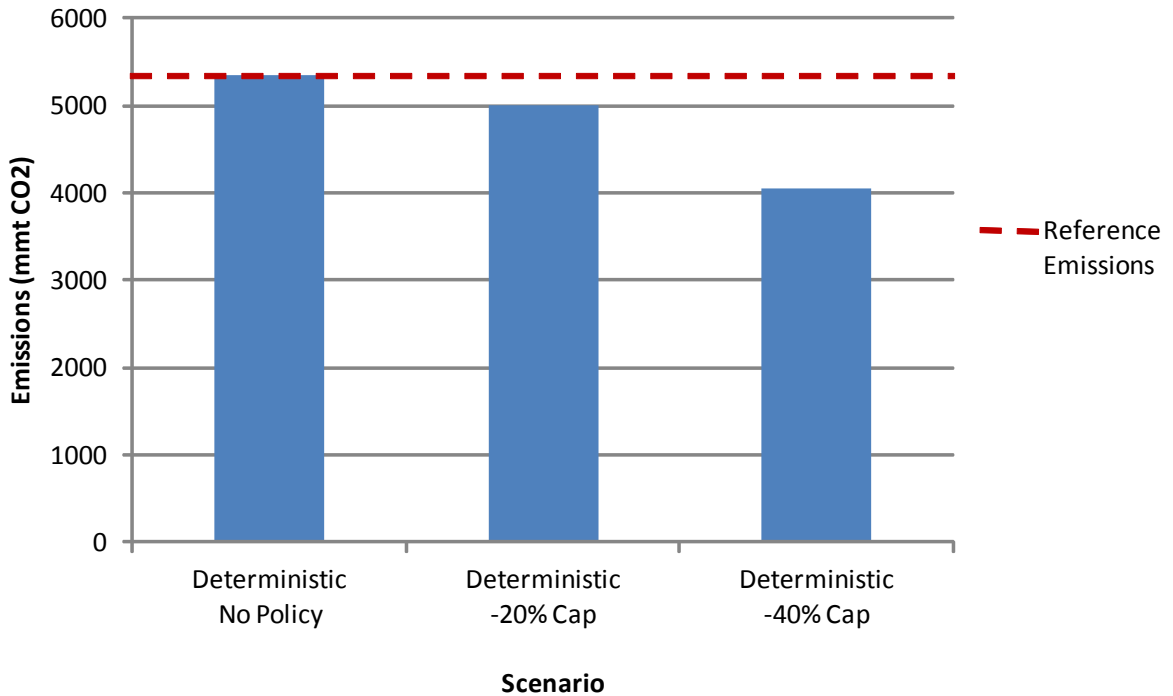
Figure 5.2 Stage 1 Emissions under Policy Certainty

Because the policy is known, we also know the optimal Stage 2 decision in terms of emissions reductions. Figure 5.3 shows the percentage of the total cumulative emissions reductions required for each cap that occurs in Stage 1 and the percentage that occurs in Stage 2. For both cap policies, the majority of the required reductions take place in Stage 2. However, the more stringent the policy, the more reductions take place in Stage 1 to help spread the burden of the policy over time. This general result has been shown by other studies focusing on the intertemporally optimal emissions path (e.g. Manne & Richels 1995a; Wigley *et al.*, 1996; Bosetti *et al.*, 2009). These studies show that pathways involving modest reductions in early years followed by sharper reductions later on are most cost-effective; and that the more stringent the emission reduction goal, the more near-term reductions are optimal. Ultimately, some emissions reductions now in anticipation of a future limit on emissions reduces the expected discounted cost of meeting future policy.
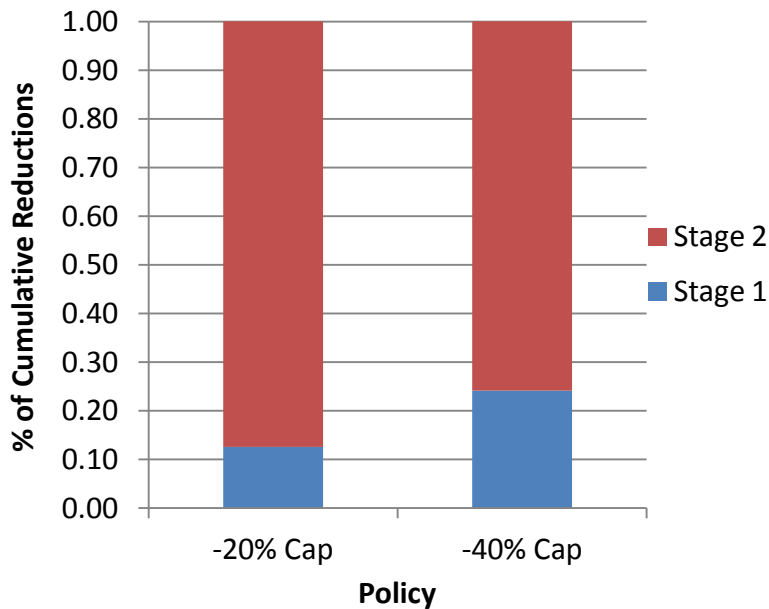
Figure 5.3 Percent of Required Cumulative Reduction Achieved in Each Stage under Certainty

5.3 Results under Uncertainty in Policy

Here, we consider uncertainty in the policy. The decision problem modeled is to choose the Stage 1 low-carbon generation share and emissions reductions with uncertainty about what the cumulative emissions limit will be in Stage 2. Because there is no objective measure of policy uncertainty, the probabilities assigned to future policy options are expectations about what future policy will be. In this modeling context, the expectations are those of a central planner and decisions are made using those expectations.

We initially consider four uncertainty scenarios, each defined by a probability distribution over the three emissions policies (Figure 5.4). The scenario nomenclature in Figure 5.4, and future figures, follows the format: .XX No Policy = probability of no policy, .XX -20% Cap = probability of -20% cap, and .XX -40% = probability of -40% cap. The results of interest are the Stage 1 decisions; Stage 2 decisions will depend on the realized policy, and therefore there are three different possible sets of optimal decisions for Stage 2, conditional on the emissions policy that is revealed. The explicit consideration of uncertainty in policy results in a near-term hedging strategy in terms of the amount of low-carbon generation in the electricity mix.

In the first uncertainty scenario, in which each policy has a 1/3 probability, the optimal Stage 1 mix of new electricity investment consists of 20% low-carbon, 63% natural gas, and 17% coal. This mix is distinct from the optimal mix resulting from any of the three deterministic scenarios. Nor is this mix the average or a linear combination of the results from the three deterministic scenarios. The resulting strategy is the solution to the dynamic programming problem as framed, and can only be determined with uncertainty analysis. The Stage 1 optimal strategy is a function of the probability distribution over the policies. In the second uncertain scenario in which either a 20% cap or a 40% cap is most likely, the optimal Stage 1 low-carbon share is 25%. A 25% share of low-carbon generation in electricity is also optimal in the fourth uncertain scenario where either no policy or a 40% cap is most likely. In that case the higher probability of a stringent target is enough to justify more near-term investment in the low-carbon technology. On the other hand, when no policy or a 20% cap is most likely (third uncertain scenario) only a 5% share of low-carbon generation is optimal in Stage 1. These scenarios demonstrate that optimal Stage 1 decisions change when uncertainty in expectations about future policy is considered, and that the optimal strategy varies with the probabilities of the policies.

100

Overall, the optimal strategy under uncertainty is to invest earlier in low-carbon technologies before the long-run emissions target is known.
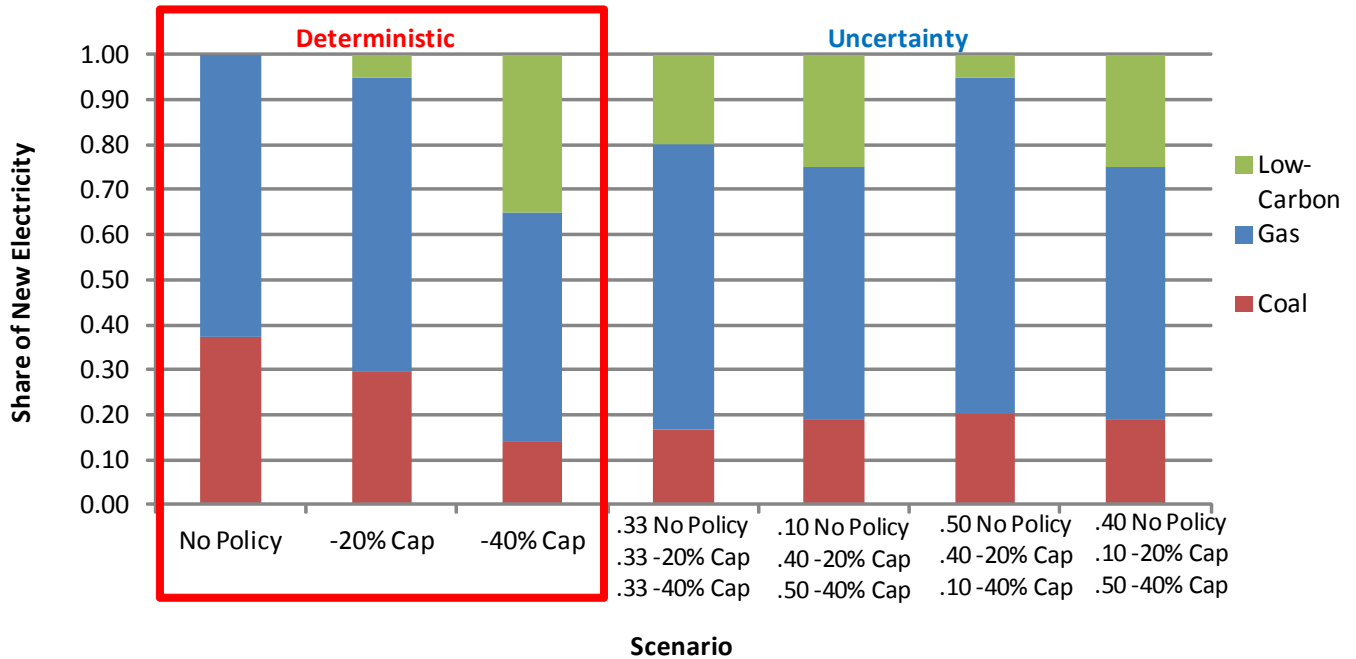
Figure 5.4 Stage 1 Shares of New Electricity under Policy Uncertainty

The Stage 1 decision about emissions reductions is also dependent on the probability distribution of future policies. Under uncertainty, the optimal strategy is to reduce emissions more than in the deterministic 20% case but not as much as in the deterministic 40% case (Figure 5.5). In the scenario that places most weight on no policy or a moderate cap, it is optimal to reduce Stage 1 cumulative emissions by 13% relative to reference. In the three other uncertainty scenarios, an 18% reduction in Stage 1 is optimal. Although those three scenarios have the same emissions reduction, only two of the three have the same optimal electricity mix (Figure 5.4). The reductions under uncertainty differ from what is optimal under the scenarios where the policy is known for certain. Thus there is a hedging strategy in terms of emissions reductions as well as in terms of low-carbon investment. Here the general strategy is to reduce emissions in Stage 1 in anticipation of policy in Stage 2. As with the electricity mix decision, the actual emission reduction hedging strategy depends on the probabilities.
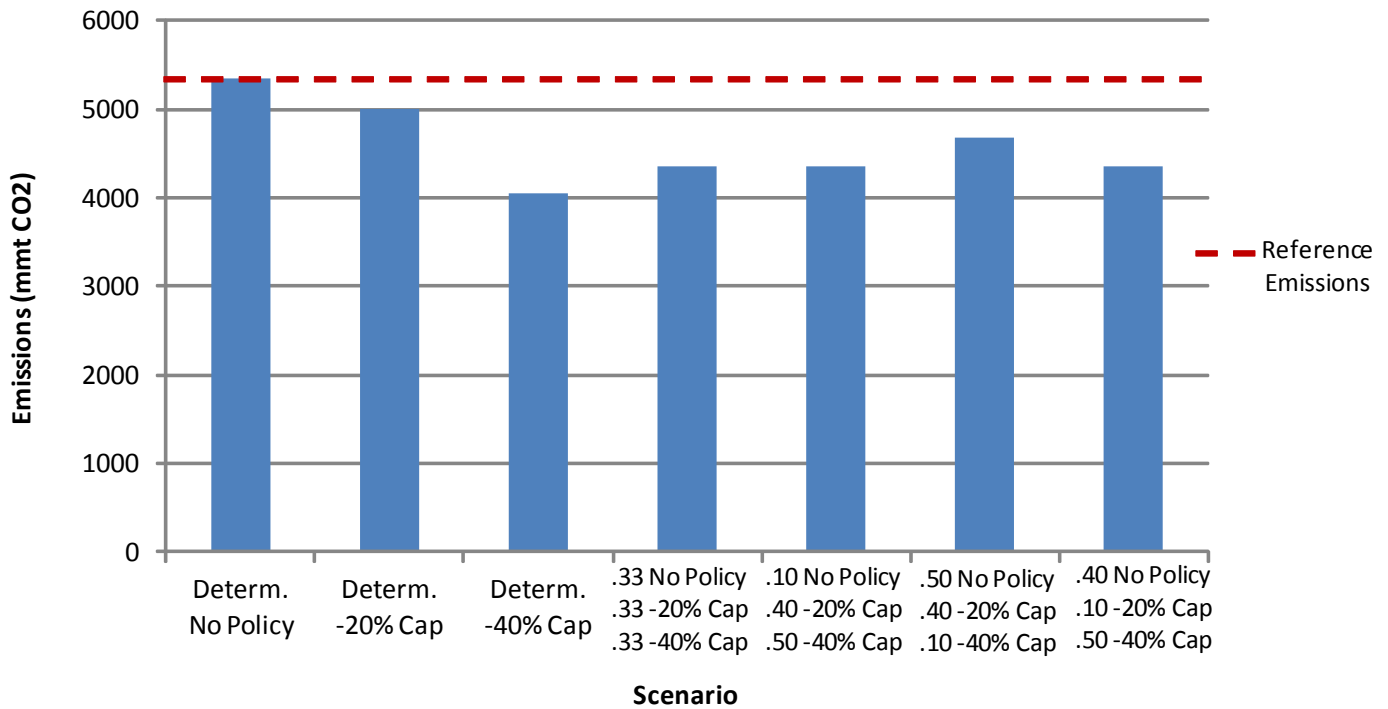
Figure 5.5 Stage 1 Emissions under Policy Uncertainty

Since it is clear that the probabilities assigned to the policies affect the optimal Stage 1 decisions, it is useful to investigate more than a handful of scenarios and explore the probability space in a systematic manner. **Error! Reference source not found.** provides a visual illustration of how the probability distributions over policies affect the optimal electricity mix. The y-axis represents the probability of no policy and the x-axis represents the probability of a 20% cap. The remaining probability $(1 - x - y)$ is therefore the probability of a 40% cap. For each probability scenario, the numbers indicate the optimal Stage 1 share of new electricity from low-carbon technologies, coal, and natural gas, as well as the optimal Stage 1 cumulative emissions reductions as percent below reference. The lower left corner of each plot is the scenario in which there is 0% probability of no policy and 0% probability of a 20% cap, and thus 100% probability (certainty) of a 40% cap. In this scenario, the optimal Stage 1 electricity mix is 35% low-carbon, 14% coal and 51% natural gas and the optimal Stage 1 emissions reduction is 24% (Figures 5.4 & 5.5). Moving away from the lower left corner, the stringent cap becomes less likely. As one would expect, greater low-carbon generation shares and more emissions reductions are optimal with higher probabilities of the 40% cap.

Share Gas in Stage 1

Probability No Policy	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
1	0.00										
0.9	0.05	0.00									
0.8	0.20	0.05	0.00								
0.7	0.20	0.20	0.05	0.05							
0.6	0.20	0.20	0.20	0.05	0.05						
0.5	0.25	0.20	0.20	0.20	0.05	0.05					
0.4	0.25	0.25	0.20	0.20	0.20	0.05	0.05				
0.3	0.30	0.25	0.25	0.25	0.20	0.20	0.05	0.05			
0.2	0.30	0.30	0.25	0.25	0.25	0.20	0.20	0.05	0.05		
0.1	0.35	0.35	0.35	0.25	0.25	0.25	0.20	0.20	0.05	0.05	
0	0.35	0.35	0.35	0.35	0.35	0.25	0.25	0.25	0.20	0.05	0.05

Probability -20% Cap

Probability No Policy	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
1	0.37										
0.9	0.30	0.28									
0.8	0.17	0.30	0.28								
0.7	0.17	0.17	0.20	0.30							
0.6	0.17	0.17	0.17	0.20	0.30						
0.5	0.19	0.17	0.17	0.17	0.20	0.30					
0.4	0.19	0.19	0.17	0.17	0.17	0.20	0.30				
0.3	0.12	0.19	0.19	0.19	0.17	0.17	0.20	0.30			
0.2	0.12	0.12	0.19	0.19	0.19	0.17	0.17	0.20	0.30		
0.1	0.14	0.14	0.14	0.19	0.19	0.19	0.17	0.17	0.20	0.30	
0	0.14	0.14	0.14	0.14	0.14	0.19	0.19	0.19	0.17	0.20	0.30

Probability -20% Cap

Share Gas in Stage 1

Probability No Policy	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
1	0.63										
0.9	0.65	0.72									
0.8	0.63	0.65	0.72								
0.7	0.63	0.63	0.75	0.65							
0.6	0.63	0.63	0.63	0.75	0.65						
0.5	0.56	0.63	0.63	0.63	0.75	0.65					
0.4	0.56	0.56	0.63	0.63	0.63	0.75	0.65				
0.3	0.58	0.56	0.56	0.56	0.63	0.63	0.75	0.65			
0.2	0.58	0.58	0.56	0.56	0.56	0.63	0.63	0.75	0.65		
0.1	0.51	0.51	0.51	0.56	0.56	0.56	0.63	0.63	0.75	0.65	
0	0.51	0.51	0.51	0.51	0.51	0.56	0.56	0.56	0.63	0.75	0.65

Probability -20% Cap

Emissions Reductions in Stage 1

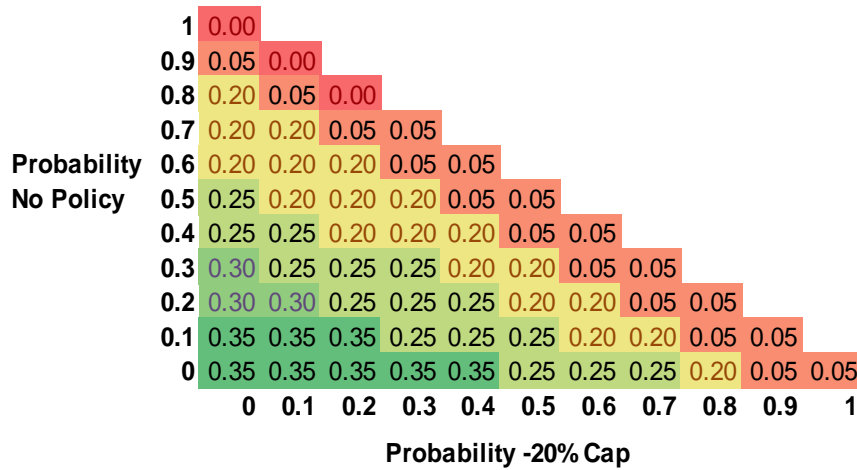| Probability No Policy	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
1	0.00										
0.9	0.06	0.06									
0.8	0.18	0.06	0.06								
0.7	0.18	0.18	0.13	0.06							
0.6	0.18	0.18	0.18	0.13	0.06						
0.5	0.18	0.18	0.18	0.18	0.13	0.06					
0.4	0.18	0.18	0.18	0.18	0.18	0.13	0.06				
0.3	0.24	0.18	0.18	0.18	0.18	0.18	0.13	0.06			
0.2	0.24	0.24	0.18	0.18	0.18	0.18	0.18	0.13	0.06		
0.1	0.24	0.24	0.24	0.18	0.18	0.18	0.18	0.18	0.13	0.06	
0	0.24	0.24	0.24	0.24	0.24	0.18	0.18	0.18	0.18	0.13	0.06

Probability -20% Cap

Figure 5.6 Optimal Stage 1 Electricity Mix and Emissions Reductions by Policy Probability

It is noteworthy is that there are very few instances in which it makes sense to use no low-carbon technology in Stage 1—only when there is no chance of a 40% cap and only a small chance (20% or less) of a 20% cap. In all other scenarios, it is best to invest in at least 5% low-carbon generation, and shift from coal to natural gas. As mentioned above, the 5% low-carbon investment is enough to overcome the fixed factor constraint so there is flexibility to greatly increase low-carbon penetration in Stage 2 if needed without facing this constraint, making it a wise hedge even with a relatively low probability of a 40% cap. If there is at least a 20% probability of a 40% cap, then it is optimal to invest in 20% low-carbon generation or more. Similarly, unless you are certain there will be no policy, it is wise to reduce Stage 1 emissions by at least 6%. If there is at least a 20% probability of a 40% cap, then it is optimal to reduce emissions by 18% or more.

The optimal shares of coal and natural gas adjust with the amount of low-carbon generation and emissions reductions, but not necessarily in a linear or obvious way. An optimal near-term strategy with less low-carbon generation generally involves more of both natural gas and coal in the mix. However there are instances when less low-carbon generation in the optimal mix is accompanied by *less* coal. For example, notice the areas of the plots that go from 25% to 20% low-carbon and correspond to 19% and 17% coal. In these cases, a lower amount of low-carbon generation is accompanied by a lower amount of coal. This is the consequence of interactions between the emissions reductions and the low-carbon penetration. Notice that the area of the plots that include 25%-20% low-carbon and 19%-17% coal, corresponds to a single 18% emissions reduction strategy. Thus we have the same reduction strategy (18%) with two different low-carbon strategies (25% and 20%). For a given emission reduction strategy, if more low-carbon generation is optimal (because of the probabilities assigned to the policies), then there does not need to be as much switching from coal to natural gas in order to reduce emissions, so more coal can be used. To achieve 18% emissions reductions without low-carbon generation would require new electrctity to be predominantly (90%) natural gas (and existing capacity would also have to shift toward natural gas). Having more low-carbon generation to achieve the emissions reductions alleviates the need for as much coal to gas shifting. As has been shown in other work (e.g. Morris *et al.*, 2010), it also decreases the shadow price on emissions, making fossil fuels less expensive. So for a given reduction target, more low-carbon generation in the optimal mix allows for more coal, and less low-carbon generation allows for less coal.

This same effect is seen in the areas of the plots corresponding to 35%-30% low-carbon, 14%-12% coal, and 24% emissions reductions.

5.4 Survey of Expectations about Future Policy

It is clear that the answer to the question of which technologies to invest in when the policy is uncertain is dependent on the expectations of different policies, which are difficult to ascertain for the real world. Individual investors have their own perceptions as to the probabilities of different future policies and are using those expectations to guide their investment decisions. These different investor expectations and decisions results in the overall electricity mix.

To get a sense of what expectations real world investors might use to make their decisions, a survey of industry experts in the area of electricity was conducted. The survey was designed to provide an initial identification of the type of expectations present among industry participants. The survey (see Appendix B) asked participants to assign probabilities to different cap policies. The online survey was sent via email to 15 individuals representing 7 different organizations, including investor-owned electric utility companies and research organizations, under the condition that individual respondents and the identity of the organizations would not be revealed. 7 of the 15 contacted responded to the survey. While more than three policy options were provided in the survey, results were later categorized into the three policy scenarios focused on in this work (see Appendix B). The results are provided in Table 5.1.

Interestingly, the respondents reflect a divided view of expectations about policy, with roughly half placing the vast majority of probability on no policy or a moderate 20% cap (responses highlighted in blue in Table 5.1) and the other half placing most probability on a moderate 20% cap or stringent 40% cap (responses highlighted in green in Table 5.1). The average expectation of the "weak policy group" (in blue) is a 69% probability of no policy, a 26% probability of a 20% cap and a 5% probability of a 40% cap (Average 1). In contrast, the average expectation of the "strong policy group" (in green) is a 15% probability of no policy, a 42% probability of a 20% cap and a 43% probability of a 40% cap (Average 2). These expectations result in very different investment decisions. According to the DP-CGE model, the "weak policy group" would pursue 5% low-carbon and 6% emissions reductions in Stage 1 while the "strong policy group" would pursue 25% low-carbon and 18% emissions reductions. The

new electricity investment mix that would result would be a combination of the decisions made under these two very different sets of expectations. That mix is likely different than what would result by using the aggregate average expectations across all respondents (Average 3).

Table 5.1 Survey Results: Expected Probabilities of Future Policies

		Expected Probability of Each Policy (%)		
		No Cap	20% Cap	40% Cap
Survey Respondent	1	77.5	22.5	0.0
	2	75.0	25.0	0.0
	3	75.0	25.0	0.0
	4	47.5	32.5	20.0
	5	20.0	50.0	30.0
	6	17.5	42.5	40.0
	7	7.5	32.5	60.0
Average 1		68.8	26.3	5.0
Average 2		15.0	41.7	43.3
Average 3		45.7	33.3	21.0

While the exact causes of differing expectations about policy are beyond the scope of this work, this survey suggests that individuals within the electric power sector may fall into one of two camps: those that believe aggressive government policy is likely and those that do not. In general, there are two main drivers of policy uncertainty: political indecision and scientific uncertainty. While ever-changing and divergent politics are partly responsible for policy uncertainty, so too is the uncertainty in the science of climate change, which makes it difficult to determine the level of emissions reduction action that should be taken. Those within the "weak policy group" may expect political gridlock, future political leadership that is not in favor of climate policy, unresolved climate science that prevents action, or climate science suggesting that climate change is not a severe problem. On the other hand, those within the "strong policy group" may expect future political leadership that is in favor of climate policy or climate science demonstrating that climate change is a severe problem that requires aggressive action to address.

Regardless of the causes, this divided view has implications for policy and modeling. It indicates the difficulty of building consensus on future policy. Different investment strategies today mean that companies that guess correctly about the future policy environment will profit or

be a better position in the future, while those who guessed wrong will lose or be worse off in the future. Recognizing this, companies would have an incentive to try to affect the future policy outcome to avoid being on the losing side. These differing views and goals may result in a political economy in which consensus on climate policy is difficult. This is problematic because, as the following section shows, policy uncertainty can be costly, meaning that consensus is valuable.

Further, from a CGE modeling perspective this divided view suggests the potential value of representing multiple representative agents with different expectations and utility functions in future work. This would enable study of how policy impacts the different agents differently, as well as how the investments made by those with divergent expectations differ from those made from the perspective of a central planner.

5.5 Effects of Uncertainty

5.5.1 The Cost of Uncertainty in the Policy

Using the DP-CGE model we can estimate the economy-wide cost of policies as well as the cost of uncertainty in the policy (the cost borne as a direct result of the policy uncertainty, which is additional to the cost of meeting the cap). The top two sections of Table 5.2 show the policy costs in terms of change in cumulative economy-wide consumption relative to the reference no policy case. The cost is shown both as a percentage change and in billions of dollars. The top section shows these costs for the scenarios in which in Stage 1 it is known with certainty what the Stage 2 policy will be (i.e. there is perfect information). In that case, the 20% cap results in -0.08% loss or $73 billion, and the 40% cap results in -0.22% loss or $205 billion. We can then calculate the expected value using a perfect prediction (EVPP) by assigning probabilities to which policy will be known for certain and taking the expected value (see Equation 5.1). For example, assume a 33% chance that in Stage 1 we will know for certain there will be no policy, a 33% chance we will know for certain there will be a 20% cap, and a 33% chance we will know for certain there will be a 40% cap (i.e. $p1=p2=p3=1/3$). In this case, EVPP of the policy cost is -0.10% or $93 billion.

$$EVPP = p1*(\text{cost when certain no policy}) + p2*(\text{cost when certain 20\% cap}) \quad \text{(EQ. 5.1)}$$
$$+ p3*(\text{cost when certain 40\% cap})$$

Table 5.2 Policy Costs: Certainty vs. Uncertainty and the Cost of Uncertainty

(1) Deterministic Policy Cost
Stage 1 Strategy Pursued Knowing Stage 2Policy for Certain
(Change in Consumption Relative to Reference)

Stage 2 Policy	% Change	Change in $billions
No Policy	0.00%	0
-20% Cap	-0.08%	-73
-40% Cap	-0.22%	-205
EVPP	-0.10%	-93

(2) Policy Cost with Uncertainty (1/3 probability for each policy)
Hedging Strategy Pursued in Stage 1, then learn Stage 2 Policy
(Change in Consumption Relative to Reference)

Stage 2 Policy	% Change	Change in $billions	
No Policy	-0.08%	-72	
-20% Cap	-0.11%	-104	
-40% Cap	-0.25%	-227	
EVUU	-0.15%	-134	

(3) Cost of Uncertainty
(Change in Policy Cost: Uncertainty vs. Deterministic)

Stage 2 Policy	Change in $billions
No Policy	-72
-20% Cap	-32
-40% Cap	-22
EVPI	42

Section 2 of Table 5.2 shows the policy costs for the scenarios in which policy is uncertain in Stage 1 and each policy is assigned a 1/3 probability, then in Stage 2 the policy is revealed to be either no policy, a 20% cap or a 40% cap. Under uncertainty a single Stage 1 hedge strategy decision is made—when each policy is assigned a 1/3 probability, the optimal hedge is 20% low-carbon, 63% natural gas, 17% coal and an 18% reduction in emissions. Three different sets of Stage 2 decisions are made depending on the policy that is ultimately revealed. Following the optimal hedge strategy in Stage 1 and then the policy-specific optimal strategy in

Stage 2 results in three different policy costs. If there turns out to be no policy, the consumption loss relative to reference is -0.08% or $72 billion. If the policy ends up being a 20% cap, the policy cost is -0.11% or $104 billion. If the policy turns out to be a 40% cap, the policy cost is -0.25% or $227 billion. The expected value under uncertainty (EVUU) is equivalent to the Bellman value—Stage 1 consumption from following the optimal hedge strategy plus the expected value of Stage 2 consumption. Comparing that total value to the total value from the deterministic reference case gives an EVUU of the policy cost of -0.15% or $134 billion.

Section 3 of Table 5.2 uses the above results to calculate the expected value of perfect information (EVPI). EVPI is a common metric used in uncertainty studies that reflects the price the decision maker would be willing to pay in order to gain access to perfect information. It is defined as the difference between EVPP and EVUU. The inverse of EVPI is the expected cost of uncertainty- the additional policy cost expected to be borne as a direct result of uncertainty in the policy. We can think of this as the cost of political indecision or of missing information. If the cumulative emissions limit was chosen and known in advance, optimal investment and emissions decisions could be made anticipating that policy. When the policy is uncertain, the best we can do in Stage 1 is to pursue the optimal hedging strategy. In this case, EVPI, and hence the cost of policy uncertainty, is $42 billion. Uncertainty increases the expected cost of policy by over 45%.

Section 3 of Table 5.2 also shows the added cost of uncertainty for each of the three policies, comparing consumption loss in the scenarios with uncertainty to the deterministic scenarios. These results drive the EVPI. If there is no policy, uncertainty costs an additional $72 billion. This is because the optimal hedge strategy pursued expensive low-carbon technology and emission reduction investments that turned out to be unnecessary in the absence of policy. If there is a 20% cap, uncertainty costs an additional $32 billion, reducing consumption by 43% compared to when the policy is certain. In this case, the hedging strategy also overinvested in low-carbon generation and emissions reductions relative to what was required to meet the 20% cap. If it was known ahead of time that the policy would be a 20% cap, it would have been best to pursue 5% low-carbon generation (instead of 20%) and 6% reductions (instead of 18%) in Stage 1. If there is a 40% cap, uncertainty costs an additional $22 billion, reducing consumption by 11% compared to the policy with certainty. In this case the hedging strategy underinvested in low-carbon generation and emissions reductions and overinvested in conventional technologies. If it was known ahead of time that the policy would be a 40% cap, it would have been best to

pursue 35% low-carbon generation (instead of 20%), 51% natural gas (instead of 63%), 14% coal (instead of 17%) and 24% reductions (instead of 18%) in Stage 1. The cost of the uncertainty in this case is mainly driven by the overinvestment in conventional generation capacity which cannot be fully utilized in Stage 2 due to the stringent emissions constraint. It is very expensive to leave existing (vintage) capacity unused or underutilized.

These results demonstrate that uncertainty in the policy is a real added cost—increasing the expected cost of policy by over 45%. Even while pursuing the optimal hedging strategy under uncertainty, once the policy is known, those hedging decisions are not necessarily optimal in retrospect. This suggests the value of setting clear, long-term policies so that decisions can be made with more complete information.

The results above also suggest a cost asymmetry between overinvesting in conventional generation and overinvesting in low-carbon generation. To investigate this asymmetry further, we consider two extreme cases: (1) it is assumed with certainty that there will be no policy, but there turns out to be a 40% cap, and (2) it is assumed with certainty that there will be a 40% cap, but there turns out to be no policy. In the first case, the policy cost of meeting the 40% cap is very high—2.3%. This high cost is due to an overinvestment in conventional generation. All Stage 1 investment is in conventional generation, and then in Stage 2, in order to meet the policy, new investment must be 50% low-carbon generation and almost 60% of existing vintage conventional generation must be left unused (stranded). The amount of stranded vintage capital is calculated by tracking in the model the total amount of generation capacity of each type available and the total amount actually used in electricity generation, the difference is available capacity that is not used (i.e. left stranded). Leaving 60% of available conventional generation capital unused indicates poor near-term investment decisions.

In the second case, the consumption loss for overinvesting in low-carbon generation is 0.13%. Although the low-carbon generation turned out to be unnecessary in the absence of policy, none of the vintage low-carbon capacity goes unused. This asymmetry is driven by the variable costs of operating. With an emissions limit, the carbon price increases the fuel cost component of conventional (fossil) generation. Even though the capital investment is sunk, the variable cost of operating the conventional generation (mainly the fuel cost) is greater than the full cost of investing in new low-carbon generation. As a result, vintage conventional generation capacity is left unused. On the other hand, low-carbon generation has low variable costs (and no

110

fuel costs), so once the capital investment is made, operation is relatively inexpensive, and lower than the full cost of investing in new conventional generation. As a result, vintage low-carbon capacity continues to be used even when there is no policy. This cost asymmetry suggests that in making investment decisions it may be wise to err on the side of too much low-carbon generation instead of too much conventional generation.

Magnitude of Policy Costs

It is worthwhile to comment on the magnitude of the policy costs. Although numbers like -0.08% and -0.25% may seem very small, the corresponding dollar amounts reveal that the policy costs are in fact very significant ($72-227 billion). More importantly, an increase in policy costs of over 45% due to uncertainty is quite significant.

To provide some context for the costs in Table 5.2, we can look to the analysis of the Waxman-Markey cap-and-trade bill (Paltsev *et al.* 2009). In that study, the medium offset case resulted in a 41% reduction in cumulative electricity emissions and a cumulative consumption loss of -1.11% or $575 billion. As expected, that policy is more costly because it is an economy-wide cap whereas the policy in this work is applied only to the electricity sector. A cap that is only applied to a subset of emissions is less stringent than an economy-wide cap, and therefore less costly (e.g. Paltsev *et al.* 2008). Given that emission reductions from the electricity sector are relatively inexpensive compared to reductions in other sectors of the economy (like transportation), it is not surprising that an electricity sector policy would cost less than half of an economy-wide policy.

Leakage of emissions is another result of a sectoral policy, which also helps reduce policy costs. Leakage refers to an increase in emissions in sectors of the economy that are not covered by the cap relative to what those emissions would have been without any policy. When only the electricity sector is subject to the cap, emissions from other sectors are allowed to increase and may grow more than they would have absent any policy. The amount of leakage largely depends on the ability to shift dirtier energy sources to uncovered sectors of the economy. Since the cap is applied to electricity, the electricity sector will need to use less fossil fuel energy, particularly coal. This may cause the price of coal to decrease. Other sectors of the economy that are able to substitute toward using more coal will do so in response to the lower coal price, and as a result these sectors will start producing more emissions. The ability of other

sectors to take advantage of cheaper conventional energy helps to offset the cost of the sectoral policy. Leakage of emissions indicates this ability to offset the policy cost through substitution. Table 5.3 shows leakage for the three deterministic scenarios as well as four uncertainty scenarios. The more stringent the policy, the more leakage occurs. Under uncertainty, the expected leakage is based on the perceived probabilities of the policies. The ability to track leakage is one of the advantages of the CGE framework which captures how all sectors are affected by a policy, not just the sector of interest.

Table 5.3 Leakage of Emissions

Scenario	Expected Leakage
Deterministic No Policy	0%
Deterministic -20% Cap	5%
Deterministic -40% Cap	10%
0.33 No Policy, 0.33 -20% Cap, 0.33 -40% Cap	7%
0.10 No Policy, 0.40 -20% Cap, 0.50 -40% Cap	8%
0.50 No Policy, 0.40 -20% Cap, 0.10 -40% Cap	5%
0.40 No Policy, 0.10 -20% Cap, 0.50 -40% Cap	7%

5.5.2 The Value of Including Uncertainty

One of the contributions of this research is demonstrating how decision making under uncertainty can be represented in a CGE model and the value of doing so. There are three main approaches to handling uncertainty in decision support models. First, we can ignore uncertainty, as is done in the two prevailing frameworks for economic models. Economic models with myopic expectations assume nothing will change in the future, while models with perfect foresight assume we know the future for certain. Second, we can acknowledge uncertainty and assess different possible scenarios as if they are certain (using either of the economic frameworks above). We can then take an "average" or "middle of the road" approach, taking the average or

middle value of the uncertain parameter as certain. Monte Carlo analysis falls into this second category because each draw of a parameter value from a distribution is used in the underlying model as if it is known to be the true value. A third option is to formally represent decision making *under* uncertainty (by using a dynamic programming framework, for example), which makes use of the imperfect expectations we have about the future to develop hedging strategies and allows for learning and revising decisions over time. These approaches result in different near-term investment strategies. In decision-making it is common practice to use expected values or best guesses in deterministic models; however these approaches can lead to poor decisions (Savage, 2012). By explicitly including sequential decision-making under uncertainty, the stochastic dynamic framework identifies the optimal hedging strategy.

The expected value of including uncertainty (EVIU) is a metric that captures the value of representing uncertainty or, conversely, the additional cost of assuming certainty. EVIU reflects the improvement in decisions obtained from explicitly representing uncertainty in the decision-making process. To demonstrate the value of including uncertainty, consider six Stage 1 strategies: (1) the optimal strategy from the DP-CGE model assuming 1/3 probability for each policy, (2) the "average" strategy- the expected value of the uncertain cap level (i.e. -20%) is imposed as though it is certain, (3) the myopic strategy (i.e. no investments in low-carbon technologies or emissions reductions until the policy is known), and (4-6) the three perfect foresight strategies- one for each of the three policies when assumed to be certain. For each of these Stage 1 strategies, the best Stage 2 strategy for each of the three possible emissions limits is identified. The expected policy cost is then calculated assuming each policy is equally likely.

Table 5.4 compares the expected policy costs for the six Stage 1 strategies, and shows that the DP strategy is the best (has the lowest expected policy cost) in the face of uncertainty. The EVIU is calculated by comparing policy costs from each strategy to those from the DP strategy. The EVIU compared to the "average" strategy is $70 billion. It is the mistake of the flaw of averages to assume the average of the scenarios would make the best strategy. In this case, doing so increases the expected policy cost by over 50% compared to pursuing the optimal hedging strategy.

113

Table 5.4 Expected Policy Cost and Expected Value of Including Uncertainty (EVIU)

Stage 1 Strategy	Expected Policy Cost (Relative to Reference)		EVIU (Stage 1 Strategy vs. DP Strategy)	
	% Change	Change $billions	% Change in Policy Cost	Change in Policy Cost $billion
Dynamic Programming	-0.15%	-134		
"Average"	-0.22%	-204	-52%	-70
Myopic	-0.83%	-761	-466%	-627
Perfect Foresight -40% Cap	-0.16%	-150	-12%	-16
Perfect Foresight -20% Cap	-0.22%	-204	-52%	-70
Perfect Foresight No Policy	-0.83%	-761	-466%	-627

Comparing the DP strategy to the myopic and perfect foresight strategies shows that ignoring uncertainty can be even more costly (Table 5.4). Pursuing the myopic strategy increases the expected policy cost by over 400% compared to pursuing the optimal hedging strategy. The investment strategy developed under perfect foresight and how the resulting policy costs compare to the optimal hedge strategy depends on what future is assumed and how that compares to the probabilistic future in the DP framework. In the experimental design from this work, assuming perfect foresight increases the policy cost by anywhere from 12% to over 400%, depending on the future emissions limit assumed. The perfect foresight strategies that assume some level of emissions limit do better than the myopic strategy in terms of expected policy cost. A myopic strategy is akin to a "wait and see" approach—only myopically cost-effective (i.e. conventional generation) investments are made in the near-term, and different investment choices can be made once the policy is revealed. However, given the experimental design, in the face of policy uncertainty it is almost always optimal to make at least some near-term investment in low-carbon generation (see Figure 5.6)—even though it is not cost-effective in the near-term, it lowers the expected cost of meeting future policy. The perfect foresight strategies that assume a limit on emissions involve near-term investment in low-carbon generation, and therefore fare better than the myopic strategy. Even if the perfect foresight strategy pursues more low-carbon generation than ultimately required to meet the policy that is implemented, it still fares better than the myopic strategy. This relates back to the cost asymmetry discussed in Section 5.4.1: it is more costly to overinvest in conventional generation than to overinvest in low-carbon generation in the face of policy uncertainty.

114

Overall, this work demonstrates that considering sequential decision-making under uncertainty results in near-term investment strategies that best minimize expected policy costs.

5.6 Impacts of Limits on Low-Carbon Generation Growth Rates

In the above results, the rate of low-carbon generation expansion is limited by the availability of fixed factor resource, as discussed in Chapter 4. This is a critical parameter because if in Stage 2 a substantial amount of low-carbon generation is needed, then there is a strong incentive to develop the technology in Stage 1. As seen from the results, the formulation of the fixed factor expansion means that as long as ~5% of investment in Stage 1 is in low-carbon generation then low-carbon generation is not significantly limited in Stage 2. To further investigate the sensitivity of results to low-carbon generation expansion rates, an alternative formulation is developed in which the low-carbon generation share in Stage 2 is strictly limited depending on the share in Stage 1. This exogenous low-carbon generation growth rate limit overrides the fixed factor.

This additional constraint limits the rate of growth of low-carbon generation as a share of new electricity between Stage 1 and Stage 2. All of the results presented in the previous sections assumed that the share of low-carbon cannot increase by more than 50 percentage points from Stage 1 to Stage 2. So if the share was 0% in Stage 1, the most it could be in Stage 2 is 50%. If the share was 20% in Stage 1, the most it could be in Stage 2 is 70%. It is possible that there is no limit on how much the share of low-carbon grows—low-carbon could constitute 0% of new electricity in Stage 1 and 100% in Stage 2. This would reflect that *all new* generation capacity put in place from 2021-2030 is low-carbon. While theoretically possible, such a solution does not seem likely or technologically feasible. All investors would have to decide to build low-carbon capacity, an unlikely prospect. Further, engineering and operational constraints (e.g. transmission constraints, reliability issues, etc.) would have to be overcome in a very short period of time in order for the electricity system to handle such large low-carbon capacity additions. However, in the past we have seen rather rapid expansion of nuclear electricity, and currently natural gas generation is quickly expanding due to the success of shale gas driving down fuel prices, suggesting there may not be much of a limit to the rate of low-carbon growth. Because it is difficult to assess and people have very different opinions about what type of low-carbon growth

rate is realistic from an engineering and technological standpoint, this section conducts sensitivity analysis on the limit to low-carbon growth rates.

Figure 5.7 illustrates the impact of the low-carbon growth rate limit assumption on the optimal Stage 1 new electricity mix when there is a 1/3 probability of each policy. If there is no limit on how much the share of low-carbon can grow, the optimal Stage 1 strategy involves 5% low-carbon. If the policy ends up being a 40% cap, the optimal Stage 2 low-carbon share is then 80%. That is a 75% increase in low-carbon's share of new electricity. In terms of share of *total* electricity (not just new electricity), low-carbon in this case increases 41 percentage points (from under 3% to almost 44%) in just ten years, reflecting a drastic change to the electricity sector in a very short period of time (particularly considering total electricity demand is growing over time). A change like this is questionable from a practical engineering standpoint. The base assumption that the share of low-carbon cannot increase by more than 50% between time periods is still optimistic. When the Stage 2 policy is a 40% cap, the optimal Stage 1 and Stage 2 low-carbon shares of new electricity are 20% and 70%, allowing the low-carbon share of total electricity to increase 34 percentage points (from about 10% to 44%). With a more pessimistic assumption that the share of low-carbon cannot increase by more than 20% between time periods, the optimal low-carbon share of new electricity is 30% in Stage 1 and 50% in Stage 2 when the Stage 2 policy is a 40% cap. This allows the low-carbon share of total electricity to increase 23 percentage points (from 14% to 37%). Overall, the more strict the limit we assume on low-carbon growth rate, the more low-carbon should be built in Stage 1. This is because higher Stage 1 low-carbon shares allow for higher Stage 2 low-carbon shares, which may be necessary to meet the stringent policy. Under strict low-carbon growth rate limits, investing in more low-carbon in Stage 1 is a way of maintaining the flexibility to achieve high low-carbon shares in Stage 2 if needed to meet the emissions limit. Ultimately, the optimal Stage 1 new electricity mix depends on what one believes is feasible in terms of low-carbon growth rates, in addition to the probabilities assigned to the policies. If you think there are strict limits to low-carbon growth rates and a non-negligible probability of a stringent policy, more Stage 1 low-carbon is optimal. If you think there are no or minimal limits on low-carbon growth rates, less low-carbon is optimal in Stage 1 even with a significant probability of a stringent emissions limit because Stage 2 low-carbon has more flexibility to grow to meet the cap.

.33 No Policy, .33 -20% Cap, .33 -40% Cap

Figure 5.7 Stage 1 Shares of New Electricity under Policy Uncertainty (1/3 Probability Each Policy) with Different Limits on Low-Carbon Growth Rate

Figures 5.8-5.10 demonstrate the effect of different low-carbon growth rate limits when there are different perceived probabilities of the policies. Appendix C includes plots that explore the full probability space of the three policies under different low-carbon growth rate limits. In general we see the pattern across probability scenarios that more probability on the stringent or moderate policy calls for more low-carbon investment in Stage 1, while more weight on no or moderate policy calls for less. Within each of the probability scenarios, the more strict the low-carbon growth rate limit assumption, the more low-carbon is optimal in Stage 1. If there is no limit to low-carbon growth rates, optimal Stage 1 low-carbon is 5% regardless of the probabilities of the policies. In this case, the 5% overcomes the fixed factor constraint and Stage 2 low-carbon can expand however much it needs to depending on the policy put in place. If there is only a 10% probability of a 40% cap, there needs to be a strict low-carbon growth rate limit (20% or 30%) for optimal Stage 1 low-carbon to be more than 5%. Alternatively, if there is only a 10% probability of no policy, unless there is no low-carbon growth rate limit, optimal Stage 1 low-carbon should be at least 20%. Overall, the Stage 1 low-carbon decision depends on both the perceived probabilities of the policies and the perceived limit on low-carbon growth rates.
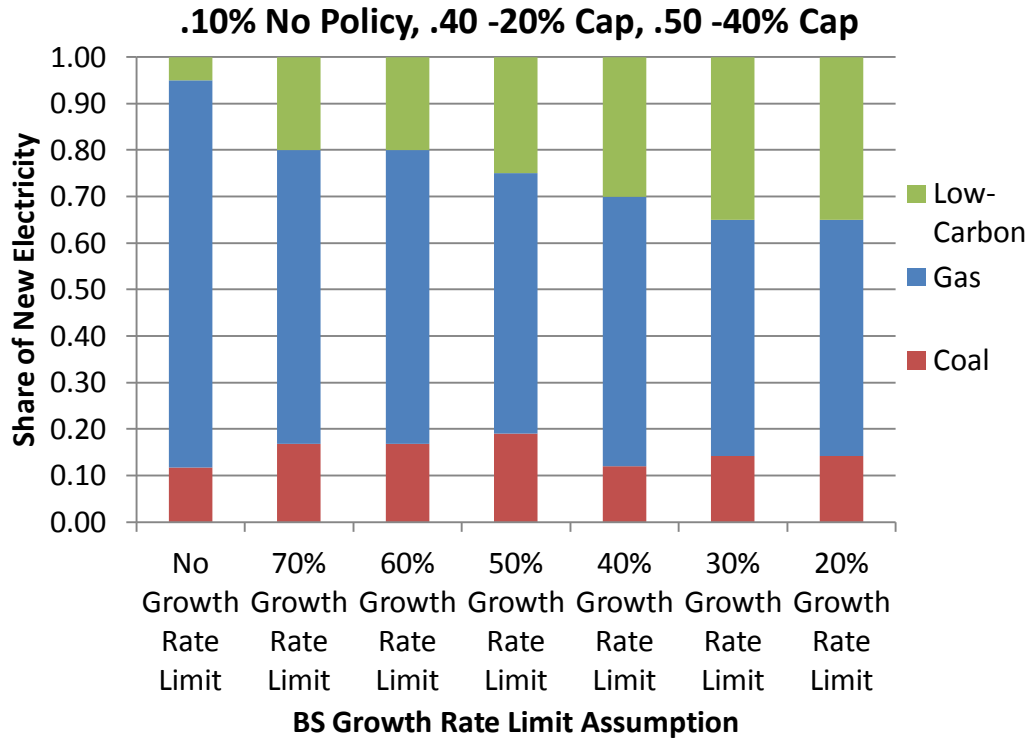
Figure 5.8 Stage 1 Shares of New Electricity under Policy Uncertainty (.10 No Policy, .40 -20% Cap, .50 -40% Cap) with Different Limits on Low-Carbon Growth Rate

Figure 5.9 Stage 1 Shares of New Electricity under Policy Uncertainty (.50 No Policy, .40 -20% Cap, .10 -40% Cap) with Different Limits on Low-Carbon Growth Rate

Figure 5.10 Stage 1 Shares of New Electricity under Policy Uncertainty (.40 No Policy, .10 - 20% Cap, .50 -40% Cap) with Different Limits on Low-carbon Growth Rate

5.7 Discussion

The results from this chapter provide insight to policy makers and those making investment decisions. There are a few approaches to thinking about what to do today: (1) we can ignore the information we have about future policy and act as though nothing will change, (2) we can pretend we know what future policy will be and hope we are right, (3) we can think through different policy scenarios, pretending they are known for certain, and then take the average or middle value of the uncertain parameter as certain, or (4) we can make use of the imperfect information we have, formally consider the uncertainty in decision-making, and develop a smart hedging strategy. The results above show the value of the smart hedge approach, which is possible using a stochastic dynamic programming framework. Other strategies increase the expected policy cost relative to the optimal hedging strategy. Given the experimental design, the expected policy cost increases by over 50% by pursuing an "average" strategy, by over 400% by pursuing a myopic strategy, and by 12% to over 400% by pursuing a perfect foresight strategy (depending on the future policy assumed).

Once we start incorporating uncertainty into decision-making, we can start thinking about strategies that minimize risk. For this problem, we want to develop near-term strategies that put us in the best position to respond to future policy in a way that minimizes the expected costs of meeting future policy. That strategy will change with the probabilities we place on future policies. The more likely we perceive a stringent policy, the more we want to do today since investments made today make it easier and cheaper to meet the future policy. This work also demonstrates the value of setting clear, long-term policies since the more certain the information available, the better decisions can be made.

Even across the range of policy probabilities, this work shows that in the near-term we should be investing in low-carbon generation and emissions reductions, and probably more than one would think. According to this model, if there is at least a 20% probability of a 40% cap, then it is optimal to invest in 20% low-carbon or more and reduce emissions by 18% or more in the next ten years. Beyond the specific numbers, the general insight is that even relatively low probabilities of a stringent future policy justify relatively aggressive near-term actions in order to minimize expected policy costs. More aggressive near-term actions are also justified if one believes there is some limit to how fast low-carbon capacity can grow. If there are such limits, more low-carbon investment now allows more flexibility in the amount of low-carbon possible in the future.

Another important consideration is the asymmetry in the cost of overinvesting in conventional generation and overinvesting in low-carbon generation. It is particularly costly if the policy turns out to be stringent enough to require previously built conventional generation capacity to go unused or underutilized due to high variable (fuel) costs, thereby stranding capital. On the other hand, if there turns out to be no policy or weak policy, previously built low-carbon generation capacity (which was more expensive to build) may have been unnecessary, but at least it can continue to be used due to its low variable costs. For this reason it may be wise to err on the side of too much low-carbon generation instead of too much conventional generation.

From a policymaking perspective, these results suggest that if there are market failures that discourage the socially optimal hedging strategy (such as spillovers of knowledge gained from investments), then a policy requiring investments in low-carbon technologies, such as a renewable or clean energy standard (CES), may be wise in the near-term while we wait for decisions to be made about a cap on emissions. Such a CES standard would essentially force all

investors to pursue a hedging strategy of near-term investments, thereby reducing the expected cost of future policy. This of course assumes that policymakers have a better understanding of the probabilities of future emissions limits than investors do, that they determine the optimal hedge strategy and implement that as the CES.

There are some additional factors not considered above that could affect the general hedging strategy. One is consideration of uncertainty in the cost of low-carbon technologies. Another is the inclusion of a safety valve provision in the emissions cap policy. These considerations and their impact on the optimal hedging strategy are the focus of the next two chapters.

(This page is intentionally left blank)

Chapter 6: Model Results: Technology Cost Uncertainty

While the previous two chapters focused on policy uncertainty, this chapter uses the DP-CGE model to investigate the impact of technology cost uncertainty on near-term electricity investment and emissions reduction decisions. Specifically, the cost of the low-carbon technology in the second stage is made uncertain. Section 6.1 gives an overview of the analysis that follows. Section 6.2 presents results from the model when the low-carbon cost is uncertain. Next, Section 6.3 explores the impact of including stochastic technological learning, in which the Stage 1 investment decision can alter the probability distribution of the low-carbon cost. Sensitivity to the parameters used to represent technological learning is explored in Section 6.4. Finally, Section 6.5 provides a concluding discussion.

6.1 Introduction

Within the CGE model, the decision of which technologies to build is driven by the relative costs of the technologies. Recall from Chapter 4 that the cost of the low-carbon technology is initially set in the model by a "markup", which is the cost relative to the conventional generation against which it competes. The amount of required inputs to production is multiplied by the markup. The reference assumption for the markup is 1.5, indicating that the low-carbon is 50% more expensive than conventional electricity in the base year of the model. In the previous chapters, this cost was constant over both decision stages. To model uncertainty in the Stage 2 cost of the low-carbon technology, the markup is allowed to vary at the beginning of Stage 2 and its value made uncertain.

As discussed in Chapter 4, the DP-CGE model uses a discrete three-point probability distribution to approximate the continuous distribution by assigning probabilities to three low-carbon technology cost scenarios: (1) a markup of 1: the low-carbon technology costs the same as conventional generation (MU1), (2) a markup of 1.5: the low-carbon continues to cost 50% more than conventional generation (MU1.5), and (3) a markup of 3: the low-carbon costs triple conventional generation (MU3).[13] As described in Section 4.3.1 in Chapter 4, these cost scenarios were informed by mean, 5[th] percentile, and 95[th] percentile estimates of future

---

[13] As an example to demonstrate the markups, if conventional generation costs 10 cents per kilowatt hour (kWh), low-carbon generation either costs 10 cents/kWh, 15 cents/kWh, or 30 cents/kWh.

technology costs from the expert elicitation literature. The extended Pearson-Tukey discrete approximation method, which assigns probabilities of 0.185, 0.65, and 0.185 to the 5th, 50th and 95th percentiles of a distribution, is then employed (Keefer & Bodily, 1983). Accordingly, the base probabilities (probabilities when there is no low-carbon investment in Stage 1) of high (MU3), medium (MU1.5) and low (MU1) cost outcomes are assigned to be 0.185, 0.63, and 0.185 respectively. Additional mean-preserving probability spreads are also explored.

In Section 6.3 stochastic technological learning is introduced by having the amount of low-carbon investment in Stage 1 affect the probabilities of the low-carbon cost markup in Stage 2. This technological learning introduces another potential motivation for being more aggressive in terms of low-carbon and emissions reductions investments in Stage 1. We have seen in previous chapters that current investment in low-carbon technologies and emissions reductions reduces the expected cost of meeting future policy by spreading the burden more between the two stages and by creating a less carbon-intensive electricity mix. Another motivation for investment in low-carbon technologies now is to proactively reduce the future cost of low-carbon technology, through learning and scale effects. If there is technological learning, such that the expected future cost of the technology decreases as the amount of near-term investment in that technology increases, then near-term investments could reduce future costs, providing greater flexibility and ease in meeting future policy. Including stochastic technological learning in the model allows for exploration of whether the potential for reducing future technology costs makes near-term investments worth it in the face of uncertainty.

6.2 Results under Low-Carbon Cost Uncertainty

First we explore the impact of uncertainty in the cost of the low-carbon technology without technological learning. Four policy scenarios are considered: (1) certain -40% cap, (2) certain -20% cap, (3) certain no policy, and (4) uncertain policy in which each policy is assumed to be equally likely (1/3 probability each)[14]. For each of these four policy scenarios, eight low-carbon cost scenarios are considered. The first three scenarios assume that the low-carbon cost—the markup (MU)—is known with certainty to either be 3, 1.5, or 1. Note that the results in

---

[14] For the remainder of this dissertation, the case in which there is a 1/3 probability of each policy is used to illustrate the impact of policy uncertainty. One should bear in mind that is just one example and that different assumptions about the probability distribution will impact results, as demonstrated in Chapter 5.

Chapters 5 and 6 assumed that the markup is deterministic with a value of 1.5. The remaining five low-carbon cost scenarios assume uncertainty in the markup at different mean-preserving spreads (Table 6.1). Each of these five scenarios result in the same expected markup (e.g. preserve the mean of the distribution), but different variances. In this way, we can isolate the effect of greater uncertainty, distinct from the effects of higher or lower average costs.

Table 6.1 Uncertain Low-carbon Cost Scenarios

Uncertain Markup Scenario	Probability of Stage 2 Low-carbon Cost Markup			Expected Markup	Variance
	1	1.5	3		
MU uncert 1	0.058	0.800	0.143	1.69	0.335
MU uncert 2	0.185	0.630	0.185	1.69	0.463
MU uncert 3	0.328	0.440	0.233	1.69	0.560
MU uncert 4	0.328	0.440	0.233	1.69	0.605
MU uncert 5	0.403	0.340	0.258	1.69	0.680

In the results of this section, the cost of the low-carbon in Stage 2 affects Stage 1 decisions through two main avenues. First, the Stage 2 markup affects the costs of emissions reductions in Stage 2, which in turn affects the optimal level of action in Stage 1. If it is going to be costly to reduce emissions in Stage 2 because the cost of low-carbon generation is high, then it will be desirable to reduce more emissions in Stage 1 in order to reduce the Stage 2 emissions reduction burden. On the other hand, if the cost to reduce emissions in Stage 2 is going to be low because low-carbon generation is low cost, then fewer reductions need to be pursued in Stage 1. Second, the Stage 2 markup influences the amount of low-carbon desired in Stage 2, which may be constrained by the limit on the low-carbon growth rate from Stage 1 to Stage 2 (see Section 5.5). The reference assumption about the maximum rate of low-carbon growth between stages allows the share of low-carbon in new investment to increase by no more than 50 percentage points from Stage 1 to Stage 2 (for example is 20% of investment is in the low-carbon technology in Stage 1, then the max low-carbon share in Share 2 is 70%). If the low-carbon will be lower cost in Stage 2, more low-carbon investment in Stage 1 is desirable because it will allow a greater share in Stage 2. This increased low-carbon could lead to lower Stage 1 emissions, even though otherwise it would be desirable to emit more in Stage 1 and less in Stage 2 when it is less costly to do so.

Table 6.2 shows the optimal Stage 1 electricity investment and emissions reduction strategy under different low-carbon cost scenarios and policy scenarios. In all cases, new electricity investments are responsible for approximately 40% of all generation.

6.2.1 Deterministic Low-Carbon Technology Cost

First, let us focus on the scenarios when the low-carbon cost is known with certainty and is either higher (MU3) or lower (MU1) than the base assumption of 1.5, which was used for the results in Chapter 5. In order to meet the stringent 40% cap when the low-carbon is expensive (MU 3), it is best to rely on natural gas and to consume less electricity overall due to the high prices of electricity (since both low-carbon and convention generation is expensive). In that case, in Stage 1 it will be optimal to invest in 15% low-carbon, 78% natural gas and 7% coal generation and reduce emissions by 35% below reference. Significant Stage 1 emissions reductions are optimal in this case because emissions reductions are more costly in Stage 2 due to the high low-carbon cost. Therefore, lower Stage 1 emissions ease the burden of reducing emissions in Stage 2. When the markup is 1.5, the optimal Stage 1 decision is a new electricity investment mix of 35% low-carbon, 51% natural gas and 14% coal and emissions reductions of 24% (same result as Section 5.2). When the markup is 1, more low-carbon is optimal in Stage 1. The optimal Stage 1 strategy is 65% low-carbon, 18% natural gas and 17% coal and 16% emissions reductions. Fewer Stage 1 emissions reductions are required in this case since emissions can be easily reduced in Stage 2 by using an abundance of low cost low-carbon.

With a markup of 1, the model continues to choose an investment mix instead of investing in all conventional or all low-carbon generation. There are several reasons for this. The actual relative cost of the technologies changes over time as the input prices change. For example, as natural gas or coal prices increase, the cost of conventional generation increases, changing the relative cost of the technologies. Also, low-carbon generation includes an initial adjust cost, and when that adjustment cost is overcome (by investing in sufficient low-carbon capacity), the relative cost of the two technologies will change. While the markup reflects a starting point for the relative cost of the technologies, there are a number of important endogenous dynamics represented that change that relative cost over time.

Table 6.2 Optimal Stage 1 Strategies under Low-carbon Cost and Policy Scenarios

Policy Scenario	Stage 1 Decision		Certain MU3	Certain MU1.5	Certain MU1	MU uncert 1	MU uncert 2	MU uncert 3	MU uncert 4	MU uncert 5
								Markup Scenario		
Certain -40% Cap	**Share of New Investment**	**Low-Carbon**	15%	35%	65%	35%	35%	35%	35%	35%
		Gas	78%	51%	18%	51%	51%	51%	51%	51%
		Coal	7%	14%	17%	14%	14%	14%	14%	14%
	Reductions	**Emissions**	-35%	-24%	-16%	-24%	-24%	-24%	-24%	-24%
Certain -20% Cap	**Share of New Investment**	**Low-Carbon**	0%	5%	45%	5%	5%	0%	0%	0%
		Gas	82%	65%	30%	65%	65%	82%	82%	82%
		Coal	18%	30%	25%	30%	30%	18%	18%	18%
	Reductions	**Emissions**	-13%	-6%	-9%	-6%	-6%	-13%	-13%	-13%
Certain No Policy	**Share of New Investment**	**Low-Carbon**	0%	0%	45%	0%	0%	0%	0%	0%
		Gas	63%	63%	30%	63%	63%	63%	63%	63%
		Coal	37%	37%	25%	37%	37%	37%	37%	37%
	Reductions	**Emissions**	0%	0%	-9%	0%	0%	0%	0%	0%
Policy Uncertainty (1/3 Probability Each Policy)	**Share of New Investment**	**Low-Carbon**	5%	20%	55%	20%	20%	15%	15%	15%
		Gas	7%	17%	21%	17%	17%	15%	15%	15%
		Coal	88%	63%	24%	63%	63%	70%	70%	70%
	Reductions	**Emissions**	-24%	-18%	-11%	-18%	-18%	-18%	-18%	-18%

Note: The MU uncert scenarios are defined in Table 6.1. They all have the same mean and are ordered by increasing variance (i.e. higher probabilities of the low and high cost outcomes). MU_uncert_2 uses the base Pearson-Tukey probabilities.

In the 20% cap and no cap scenarios, an MU of 3 leads to no low-carbon and an MU of 1 leads to 45% low-carbon in Stage 1. Even when there is no policy, if the markup is 1 it is optimal to make significant investments in low-carbon in Stage 1. This is because relying more on low-cost low-carbon electricity in Stage 1 frees up coal and natural gas resources for use in other sectors at a lower price (since the demand for these resources from the electricity sector decreases).

In the uncertain policy scenario, defined here as 1/3 probability of each policy, when the markup is 3 the optimal Stage 1 investment in low-carbon is 5%, in contrast to 15% under the 40% cap case and 0% under the other policy cases. Similarly, the optimal Stage 1 emissions reductions are 24%, in contrast to 35% reduction under the 40% cap and 13% reduction under the 20% cap. This hedging strategy protects against the particularly high risk associated with there being a 40% cap and a markup of 3, which would make it very expensive in Stage 2 to meet the cap (if enforced) due to costly or uneconomical low-carbon generation. The 24% emissions reduction hedge in Stage 1 helps to ease the burden of expensive emissions reductions in Stage 2 that may be required depending on the policy ultimately implemented. When the markup is 1.5, the optimal Stage 1 decision is an electricity mix of 20% low-carbon, 63% natural gas and 17% coal and emissions reductions of 18% (same result as Section 5.3). When the markup is 1, the optimal Stage 1 strategy is 55% low-carbon, 24% natural gas and 21% coal and 11% emissions reductions. In that case, Stage 1 low-carbon investment has value regardless of the policy that is ultimately implemented because it allows low-carbon investment to grow without limit in Stage 2 if it turns out to be low cost. Fewer Stage 1 emissions reductions are required in this case since emissions can be easily reduced in Stage 2 by using an abundance of low cost low-carbon.

6.2.1 Uncertain Low-Carbon Technology Cost

Next, let us consider the scenarios in which the low-carbon cost is uncertain (also in Table 6.2). If it is known with certainty there will be a 40% cap or certain there will be no cap, low-carbon cost uncertainty does not affect the optimal Stage 1 strategy. For a wide range of distributions with different variances, the strategy is the same as when the markup is 1.5 with certainty. Similarly, if it is known that there will be a 20% cap or if the policy is uncertain with 1/3 probability each, lower variance distributions (MU uncert 1, 2) result in the same strategy as

when the markup is known to be 1.5. However, the higher variance probability distributions of cost result in less low-carbon 0% (instead of 5%) in the -20% cap case and 15% (instead of 20%) in the uncertain policy case. In these cases, the probability that the markup is 3 is too high to warrant as much investment in low-carbon in Stage 1. It is a safer bet to invest in less low-carbon in Stage 1 (since you may not want to use much low-carbon in Stage 2 depending on the realized cost).

These results show very weak or no impact from uncertainty in the low-carbon cost on the Stage 1 strategy. Most cases with low-carbon cost uncertainty result in the same strategy as the strategy when a 1.5 markup is certain (i.e. the "middle" strategy). There are some instances when that is not the case and the low-carbon cost uncertainty results in less low-carbon investment in Stage 1 than the "middle" strategy, but only 5% less. Other probability distributions for the low-carbon cost markup that are not mean-preserving may impact the Stage 1 strategy, but this would be caused in part by the higher or lower expected costs. However, ultimately the effect of low-carbon cost uncertainty, in terms of the causal mechanisms discussed here is small.

6.3 Results under Low-Carbon Cost Uncertainty with Stochastic Technological Learning

6.3.1 Stochastic Technological Learning

The previous section showed that two reasons for uncertainty in the low-carbon cost to impact the Stage 1 strategy—by changing the expected costs of emissions reductions in Stage 2 and the constraints on the low-carbon technology growth rate between periods—have in practice a very weak effect given the experimental design explored. There is, however, a third mechanism that some have argued for: namely, that scale effects and learning-by-doing alter the value of near-term investment by accounting for additional benefits. In this section, we explore how the inclusion of technological learning, in a stochastic setting, influences near-term optimal investment decisions.

Specifically, we explore the following question: how does the optimal Stage 1 strategy change when the amount of low-carbon investment in Stage 1 affects the probabilities of the low-carbon cost markup in Stage 2? Studies of electricity investment decisions that include learning for technology costs are typically deterministic, most utilizing learning curves in which

a given amount of investment or production or capacity leads to a given level of cost reduction (e.g. Kypreos & Bahn, 2003; van der Zwaan *et al.*, 2002; Messer, 1997; Loulou *et al.*, 2004; Seebregts *et al.*, 1999; Morris, 2002; Mattsson & Wene, 1997; Berglund & Soderholm, 2006; Kypreos & Barreto, 2000). However, here we investigate technological learning in a stochastic framework in which a given amount of capacity investment changes the *probabilities* of future technology costs. This represents a contribution to the learning-by-doing literature.

The representation of stochastic technological learning in the mode is informed by learning-by-doing (LBD) curves (also known as experience curves), which represent how the cost of technologies declines as a function of cumulative production or capacity. For electricity generation technologies, LBD curves are often developed for categories of technologies based on cumulative installed capacity (Clarke *et al.*, 2008). LBD formulations are founded upon the concept that technology improves (e.g., costs decline) as cumulative experience with the technology increases and repetition and familiarity leads to greater efficiency. Empirical research using data on cumulative installed capacities and technology costs has been used to develop learning curves for electricity technologies (e.g., Ibenholt, 2002; Colpier & Cornland, 2002; Yeh & Rubin, 2007). Such studies can be used in this work to calibrate the technological learning parameters to reflect the empirical relationship between cumulative installed capacity and technology cost reductions.

In the LBD literature, LBD curves are often expressed as power functions, for example:

$$C_q = C_0 * q^{-b}$$ (EQ. 6.1)

where C_q is the cost per unit q, C_0 is the cost for the first unit, q is the cumulative capacity or production (experience over time) and b is a so-called experience index. The value 2^{-b} is called the progress ratio (PR). If an experience curve shows a progress ratio of 85 percent it means that cost declines by 15 percent for each doubling of cumulative capacity. Some studies use the term learning rate (LR), defined as (100-PR). Studies show that progress ratios vary significantly across technologies. For energy supply technologies, studies have shown that the progress ratio varies from 80 to more than 100 percent[15] (Neij, 1997).

---

[15] A progress ratio of over 100% reflects costs increasing despite growing capacity, and is typically explained by improvements made in areas such as performance, efficiency, safety, etc.

Here, I propose a model of stochastic technological learning in which the Stage 1 low-carbon shares affect the probabilities of the Stage 2 low-carbon cost scenarios. For the three-point discrete distributions used here, the model is parameterized so that as the amount of Stage 1 low-carbon increases, the probability of a low Stage 2 markup increases, and the probability of a high Stage 2 markup decreases. Specifically,

$$P\{MU=1\} = P_0^L + BS_1 * \pi^L \qquad \text{(EQ. 6.2)}$$

$$P\{MU=3\} = P_0^H + BS_1 * \pi^H \qquad \text{(EQ. 6.3)}$$

$$P\{MU=1.5\} = 1 - P\{MU=1\} - P\{MU=3\} \qquad \text{(EQ. 6.4)}$$

where BS_1 represents the Stage 1 low-carbon share, which corresponds to a cumulative low-carbon capacity at the end of Stage 1. For a starting distribution, using the extended Pearson-Tukey method, P_0^L and P_0^H are 0.185. The values of the technological learning parameters π^L and π^H can then be calibrated to be consistent with the LBD literature, as shown below.

Calibration requires several informational components: technology cost in 2010, capacity in 2010 and 2020, and a progress ratio. Because the model defines the cost of the low-carbon technology in terms of a markup over the conventional technology, those informational components are required for both a conventional and a low-carbon technology. For our calculations we use natural gas generation for the conventional technology and renewables for the low-carbon technology. Natural gas is assigned a progress ratio (PR) of 90% (consistent with e.g. Colpier & Cornland, 2002; McDonald & Schrattenholzer, 2001) and renewables are assigned a PR of 80% (consistent with e.g. Ibenholt, 2002; van der Zwaan & Rabl, 2003; McDonald & Schrattenholzer, 2001). 2010 capacity and projected 2020 capacity for gas and renewable generation are from EIA (2013a). 2010 costs for gas and renewables are defined as the levelized cost of electricity (LCOE), which is the price of electricity per kWh taking into account capital, operating, fuel, and other costs (see Section 2.2.4 in Chapter 2). LCOE for renewables is set such that the markup of renewables over gas is 1.69, matching the expected markup using the Pearson-Tukey probabilities. The LCOE for 2020 for gas and renewables is then calculated—

using the learning curve, not the model.[16] Based on the amount of capacity added since 2010 and the progress ratio, the 2020 LCOE is reduced compared to the 2010 LCOE.

Figure 6.1 shows the change in costs for the technologies that result from this simple learning curve calculation. Dividing the 2010 cost for renewables by the 2010 cost for natural gas yields the markup of 1.69. Dividing the 2020 projected cost for renewables by the 2020 projected cost for natural gas yields a markup of 1.61, representing a 4.5% decrease in the relative cost of the low-carbon from 2010 to 2020. More renewable capacity will be added during that time period than natural gas capacity and renewables also have a higher learning rate (20% vs. 10%), which results in the renewable cost decreasing by 6.2% and the natural gas cost decreasing by 1.6%. Typical LBD studies focus on the cost reduction of a single technology (e.g., line A in Figure 6.1), but here we are focused on the reduction in the cost of the low-carbon relative to the cost of conventional generation.

Figure 6.1 Changing Costs of Technologies Using Learning Curve

The change in the markup from 1.69 to 1.61 corresponds to renewable generation capacity increasing from 126 gigawatts (GW) to 154 GW. In the DP-CGE model, different Stage

---

[16] This learning curve approach does not explicitly account for changing input prices over time (e.g. fuel, capital, labor, etc.).

1 low-carbon investment decisions will result in different levels of cumulative low-carbon capacity at the end of Stage 1. For example, a 25% share of low-carbon in Stage 1 investment corresponds to approximately 155 GW of capacity in 2020. Since the base Pearson-Tukey probabilities and zero low-carbon investment in Stage 1 result in an expected markup of 1.69, 25% low-carbon in Stage 1 should be calibrated to result in markup probabilities that have an expected markup of 1.61. So the technological learning parameters (π^L and π^H) should be set such that the expected markup when you chose 25% low-carbon in Stage 1 is 1.61. A set of parameter values that achieve this are $\pi^L = 0.3$ and $\pi^H = 0.1$. The resulting model of stochastic technological learning is illustrated in Table 6.3 using representative Stage 1 investment decisions. Less low-carbon investment in Stage 1 leads to higher probabilities of MU3 while more low-carbon investment in Stage 1 leads to higher probabilities of MU1.

Table 6.3 Stochastic Technological Learning: Probability of Stage 2 Low-carbon Cost Markup Given Stage 1 Low-carbon Share

Stage 1 Low-Carbon Share	Stage 2 Low-carbon Cost Markup			Expected Markup
	1	1.5	3	
0%	0.185	0.630	0.185	1.69
5%	0.198	0.623	0.180	1.67
10%	0.210	0.615	0.175	1.66
15%	0.223	0.608	0.170	1.64
20%	0.235	0.600	0.165	1.63
25%	0.248	0.593	0.160	1.61
30%	0.260	0.585	0.155	1.60
35%	0.273	0.578	0.150	1.58
40%	0.285	0.570	0.145	1.57
45%	0.298	0.563	0.140	1.55
50%	0.310	0.555	0.135	1.54
55%	0.323	0.548	0.130	1.52
60%	0.335	0.540	0.125	1.51
65%	0.348	0.533	0.120	1.49
70%	0.360	0.525	0.115	1.48
75%	0.373	0.518	0.110	1.46
80%	0.385	0.510	0.105	1.45
85%	0.398	0.503	0.100	1.43
90%	0.410	0.495	0.095	1.42
95%	0.423	0.488	0.090	1.40
100%	0.435	0.48	0.085	1.39

In the DP-CGE model, Stage 1 decisions about electricity technologies and emission reductions must be made without knowing which of the three low-carbon costs will be realized in Stage 2, but rather with expectations about which costs are most likely and an understanding that the cost will be driven by near-term investments in low-carbon generation.

6.3.2 Results with Stochastic Technological Learning

Table 6.4 shows the optimal Stage 1 strategies for several policy scenarios when technological learning is modeled, along with the optimal strategies from the previous section without technological learning or cost uncertainty. In the table, the expected markup for the technological learning cases can be used to identify the effective probabilities for the markups in Table 6.3. The inclusion of technological learning does not change the Stage 1 strategy for deterministic policies of 20% cap or no cap. However, if a 40% cap is known with certainty or if the policy is uncertain with each policy equally likely, technological learning does change the strategy. For the deterministic 40% cap scenario, a 95% low-carbon share and a 19% emissions reduction are optimal in Stage 1 (vs. 35% low-carbon and 24% emissions reductions in the absence of technological learning). When the policy is uncertain, a 65% low-carbon share and a 20% emissions reduction are optimal (vs. 20% low-carbon and 18% emissions reductions in the absence of technological learning). In these cases, the incremental cost of more low-carbon investment in Stage 1 is offset by the additional benefit of reducing the probability of having a markup of 3 under a 40% cap policy, which would be very costly (if enforced).

In both the -40% cap and uncertain policy cases, the strategy with technological learning involves more low-carbon investment than when it is known for certain that the markup will be 1 (95% vs. 65% for -40% cap and 65% vs. 55% for policy uncertainty). With the stochastic technological learning, the justification for the additional Stage 1 low-carbon is to bring the expected low-carbon cost down for Stage 2. You would not make the additional investments if you were certain the cost would be low regardless of actions. In other words, if low future costs are "free", you do not need to do as much in Stage 1. However, if you can pay to increase the probability of low future costs by investing more now, it may be worth it to do so. Lower future costs provide greater flexibility and ease in meeting future policy. This flexibility is particularly valuable as the probability of a stringent emissions cap increases.

134

Table 6.4 Optimal Stage 1 Strategies with Stochastic Technological Learning for the Low-Carbon Technology

	Expected Markup (MU)	Share of New Investment			Emissions Reductions
		Low-Carbon	Coal	Gas	
Certain -40% Cap					
Certain MU3	3	15%	7%	78%	-35%
Certain MU1.5	1.5	35%	14%	51%	-24%
Certain MU1	1	65%	17%	18%	-16%
Uncert MU No Learn*	1.69	35%	14%	51%	-24%
Uncert MU Learn	1.4	95%	3%	2%	-19%
Certain -20% Cap					
Certain MU3	3	0%	18%	82%	-13%
Certain MU1.5	1.5	5%	30%	65%	-6%
Certain MU1	1	45%	25%	30%	-9%
Uncert MU No Learn*	1.69	5%	30%	65%	-6%
Uncert MU Learn	1.67	5%	30%	65%	-6%
Certain No Policy					
Certain MU3	3	0%	37%	63%	0%
Certain MU1.5	1.5	0%	37%	63%	0%
Certain MU1	1	45%	25%	30%	-9%
Uncert MU No Learn*	1.69	0%	37%	63%	0%
Uncert MU Learn	1.69	0%	37%	63%	0%
Policy Uncertainty (1/3 Each Policy)					
Certain MU3	3	5%	7%	88%	-24%
Certain MU1.5	1.5	20%	17%	63%	-18%
Certain MU1	1	55%	21%	24%	-11%
Uncert MU No Learn*	1.69	20%	17%	63%	-18%
Uncert MU Learn	1.49	65%	18%	17%	-20%

* Note: Uncert MU No Learn assumes base Pearson-Tukey probabilities and corresponds to the MU_uncert 2 case in Table 6.1.

Similarly, if technological learning is deterministic, less low-carbon investment is optimal in Stage 1. Under deterministic technological learning, near-term investments determine the actual future markup instead of the probabilities of future markups. Deterministic technological learning is modeled by taking the expected markups from stochastic technological learning formulation (Table 6.3) and assuming that markup value occurs with certainty if the necessary amount of Stage 1 investment is made. For the uncertain policy case, the optimal

decision with deterministic technological learning is 25% low-carbon investment, lower than the 65% that is optimal with stochastic technological learning. If you know exactly how your investments will reduce future costs, you will only invest the minimum amount needed to achieve the desired cost reduction. In the stochastic case, you do not know how effective your investments will be at reducing future costs, so you have incentive to invest more in order to increase the chances of a low-cost outcome.

Ultimately, whether or not more near-term investment is justified depends on expectations about future policy as well as the expected low-carbon cost distribution, which is determined by the technological learning parameters. Balancing these expectations provides the optimal hedge strategy.

6.4 Sensitivity to Stochastic Technological Learning Rate

In this section sensitivity analysis is conducted on the technological learning parameters π^L and π^H in order to explore the impact of technological learning on the results. The parameters π^L and π^H determine the probabilities of the markups that result from different Stage 1 low-carbon decisions (see equations 6.2-6.4). These two parameters determine the magnitude of the technological learning—higher values lead to larger impacts on the probabilities and larger reductions in the expected markup for a given Stage 1 low-carbon decision.

The magnitude of the technological learning can be translated to learning rates. In the learning-by-doing literature the learning rate is equal to (100 – Progress Ratio) and measures how much costs change as a function of cumulative capacity. A progress ratio of 85% corresponds to a learning rate of 15% and means that cost declines by 15% for each doubling of cumulative capacity. The base technological learning parameters used in the previous section (π^L = 0.3 and π^H = 0.1) were calibrated to learning-by-doing literature estimates of learning rates of 20% for low-carbon generation and 10% for conventional generation. In the same manner, we can identify technological learning parameters that correspond to different learning rates for the low-carbon technology (holding the conventional generation learning rate at 10%). Table 6.5

illustrates values of parameters π^L and π^H that correspond to learning rates covering the typical range in the literature.[17] The higher the learning rate, the more expected costs are reduced.

Table 6.5 Technological Learning Parameters and Learning Rates (LR)

π^L	π^H	LR
0	0	0%
0.001	0.001	5%
0.050	0.050	10%
0.100	0.100	15%
0.300	0.100	20%
0.540	0.100	25%
0.575	0.180	30%

Using these parameter values we can explore how different learning rates affect the optimal Stage 1 decisions. First, consider the scenario when policy is known for certain to be a 40% cap (Figure 6.2). The 0% learning rate is equivalent to no technological learning—the cost stays constant regardless of the low-carbon investment in Stage 1. As the learning rate increases, more low-carbon is optimal in Stage 1. This is because with higher learning rates, Stage 1 low-carbon investments are more valuable as they have larger impacts on the probabilities of future markups and result in larger reductions in the expected markup. Essentially, higher technological learning rates mean you get "more bang for your buck" of low-carbon investment in Stage 1, and therefore there is incentive to invest more in low-carbon generation. With high enough technological learning rates (20-30%), the optimal share of low-carbon in Stage 1 increases to 95%, which makes the probability of MU=1 very high (47-73%), and the probability of MU=3 very low (1.4%-9%). The same effect, though weaker, is seen when the policy is uncertain (defined here as 1/3 probability of each policy) (Figure 6.3).

For a 40% cap, the optimal Stage 1 low-carbon share differs from the strategy without technological learning at a learning rate of 10%. A learning rate of 15% leads to a different decision when the policy is uncertain. Both of these rates are lower than most estimates from the literature of learning rates for advanced low-carbon technologies, which typically fall around 75-

---

[17] There are other sets of parameter values that result in different probability distributions, but the same expected markup and therefore the same learning rates. The parameter values in Table 7.5 are therefore examples of values that correspond to a set of learning rates.

80%. At those learning rates, technological learning has a significant impact on the optimal Stage 1 low-carbon share under these two policy scenarios.

Under a 20% cap, a learning rate of at least 25% is required to change the optimal Stage 1 low-carbon share, which jumps from from 5% to 95%. With a 25% learning rate, 95% low-carbon investment in Stage 1 results in about a 70% chance of a MU=1, a 9% chance of MU=3, and an expected markup of 1.29. At these probabilities, the cost of investing in more low-carbon in Stage 1 than is necessary to meet the cap is outweighed by the potential value of a markup of 1, which would allow compliance with the 20% cap at lower policy cost than if the markup turned out to be 1.5 or 3. Even when it is certain there will be no cap, a high enough learning rate encourages Stage 1 low-carbon investment. In that case, a learning rate of 30% changes the optimal Stage 1 low-carbon share from 5% to 95%. The cost of investing in low-carbon in Stage 1 that is otherwise unnecessary is outweighed by the potential value of a markup of 1, which in this case would allow electricity to be generated from low-carbon at the same cost as conventional generation, thereby freeing up coal and gas resources for other uses in the economy at lower prices and increasing overall economic consumption and therefore social welfare.
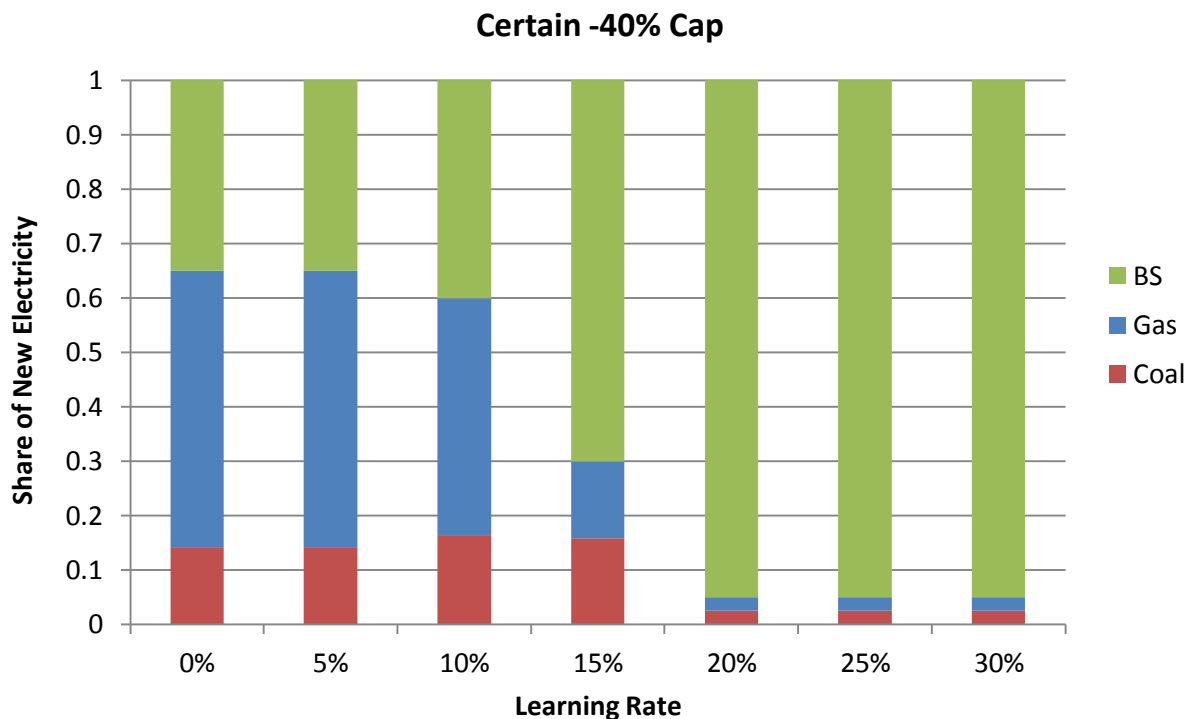
Figure 6.2 Stage 1 Shares of New Electricity under a -40% Cap with Different Learning Rates

Uncertain Policy (1/3 Probability of Each Policy)
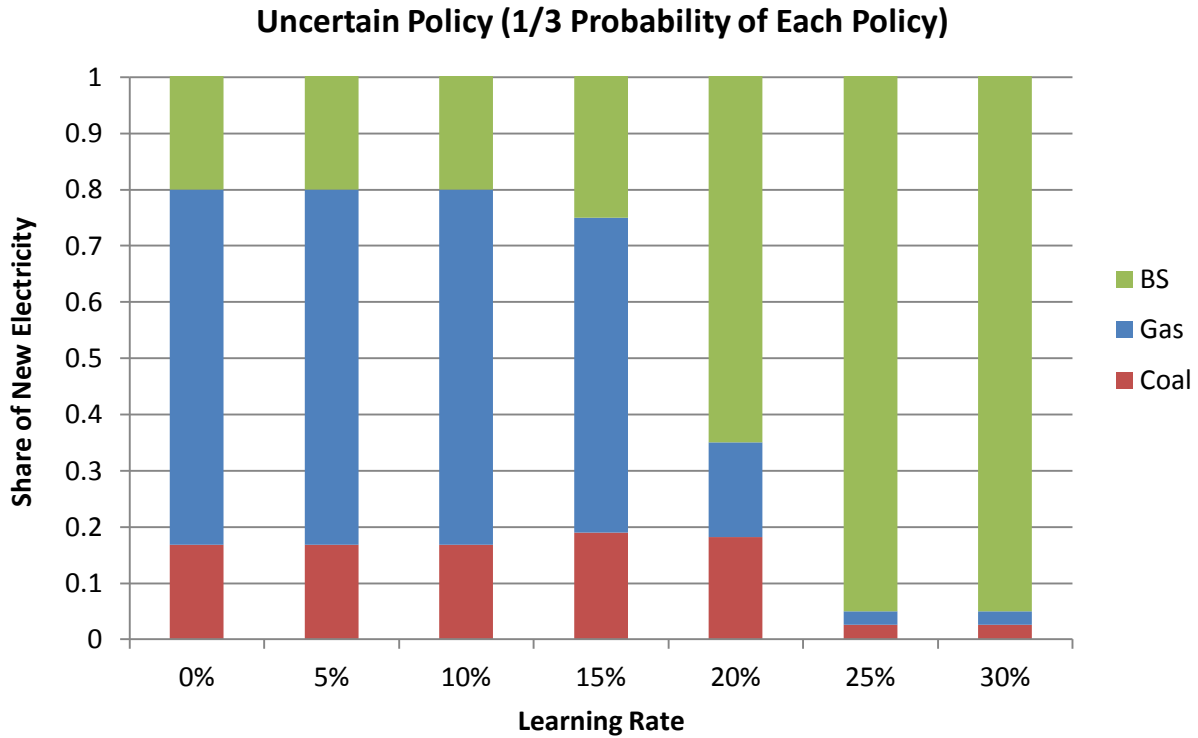
Figure 6.3 Stage 1 Shares of New Electricity under Policy Uncertainty (1/3 Probability Each Policy) with Different Learning Rates

6.5 Discussion

The results from this chapter demonstrate an additional motivation for investing in low-carbon generation in the near-term. The results from Chapter 5 show that in many cases of policy uncertainty, Stage 1 investments in low-carbon make economic sense because they lower the expected costs of emissions reductions in Stage 2. Stage 1 low-carbon investment is also optimal when there is a constraint on the low-carbon technology growth rate between periods and you expect you may need substantial low-carbon in Stage 2. *A priori*, one might expect uncertainty in the low-carbon cost to affect Stage 1 low-carbon decisions for the same two reasons, but that turned out *not* to be the case. At least given the uncertainty in the costs of low-carbon generation as represented in the literature, the cost uncertainty does not have a strong effect on State 1 strategy for the scenarios explored.

However, including stochastic technological learning in which the share of low-carbon in Stage 1 affects the probability of cost markups in Stage 2 introduces a third motivation for near-

term low-carbon investment: the ability to lower the expected cost of low-carbon in the future. Through learning and scale effects, the cost of technologies may decline with increased cumulative capacity. Depending on the rate of technological learning, as well as the expectations about future policy, the value of reducing expected future costs can provide strong motivation for additional low-carbon investment in the near-term. As the probability of a stringent cap increases, the potential need for significant amounts of low-carbon to meet the cap increases, and in turn the value of near-term investments that reduce the expected future low-carbon cost also increases. Investors should be aware of this value when making investment decisions.

These results have policy implications. In the model set up, the benefits of technological learning are fully realized and taken into account by the central planner in identifying the optimal near-term strategy. However, in a more realistic industry structure where there is competition and assuming technological learning benefits spillover to competitors or are not fully captured by the private sector investors, there may be underinvestment in low-carbon technologies compared to what is socially optimal. In this case, there may be need for government policy to encourage private investment in low-carbon technologies. Policies requiring the use of low-carbon technologies, demonstration projects, or tax incentives could be used in the near-term to encourage private investment in low-carbon technologies, and could help lower the expected cost of those technologies in the future, making it less costly to meet future policy.

From a modeling perspective, the inclusion of stochastic technological learning is a valuable step beyond the traditional learning-by-doing curves. While LBD curves capture the fact that changes in technology cost are not "free" but rather come at the expense of investments in the technology, they fail to capture the uncertainty surrounding how much investments will reduce future costs. The formulation in this chapter captures both of these effects, capturing the uncertain impact of investments on future costs by having investments affect the probability distribution of future costs. This new formulation suggests more near-term investment in low-carbon generation is optimal than would be the case if the learning effects were certain. The uncertain nature of technological learning encourages increased investments in order to lower the probability of a high cost outcome and increase the probability of a low cost outcome. This result contributes to the thinking about technological learning.

The next chapter considers the inclusion of a safety valve—a policy design feature designed to contain policy costs (e.g. in the face of technology cost uncertainty).

Chapter 7: Model Results: Policy Design with a Safety Valve

The previous chapters used the DP-CGE model to investigate near-term electricity investment and emissions reduction decisions under uncertainty in future climate policy and low-carbon technology costs. If a policy limiting emissions is implemented, the cost of meeting the policy is uncertain and partly driven by uncertainty in the cost of low-carbon technologies. The question arises as to whether an emissions cap will be enforced regardless of the cost of meeting the policy. This chapter considers how the near-term electricity investment strategy might change if future policy includes a "safety valve", which is a policy mechanism meant to contain the cost of the policy by loosening the cap. Specifically, how does knowing that the policy may be relaxed if it is too costly affect what decisions are made today? Section 7.1 reviews the safety valve mechanism and how it is implemented in this analysis. Section 7.2 presents results from the model when a safety valve is included in the policy. Section 7.3 provides a concluding discussion.

7.1 Introduction

7.1.1 A Hybrid Policy Approach: The Safety Valve

Two broad policy approaches have been generally used to address environmental issues such as carbon dioxide emission. One is to tax emissions, setting the tax equal to the marginal damage caused by the emissions (often referred to as the social cost of carbon). The other approach is to cap emissions and allow firms to trade emissions allowances. One of the characteristics of a cap-and-trade system is that the amount of emissions is known (the cap), but the carbon price is uncertain—it is determined in the market and depends on demand for allowances given available options to reduce emissions (and is therefore partially driven by uncertainty in the cost of low-carbon technologies). As a result, carbon prices can be volatile, as was the experience in the first round of the European Trading Scheme (EU ETS) beginning in 2005 (Figure 7.1). Prices have also been volatile in the U.S. Acid Rain Program. Such price volatility makes it difficult for firms to plan. Policymakers and firms particularly worry about the upside risk of a cap-and-trade system, meaning the risk of prices rising higher than expected or desired. Unexpectedly high costs hurt economic growth and could spur policymakers to loosen the cap or suspend or cancel the entire policy.

ECX Futures Contracts Settlement Prices

Dec 2008 Contracts

Dec 2007 Contracts

Source: European Climate Exchange

Figure 7.1 Permit prices in the first round of EU ETS.

(Note: The red line is futures prices.)

Some authors have argued that an emissions tax is more efficient, especially considering uncertainty (e.g. Weitzman, 1974, Pizer, 1997, Newell and Pizer, 2003). Others have favored a cap-and-trade policy because if there is a known catastrophic level of emissions, a cap can be set to certainly avoid that level. Politically, many have preferred a cap-and-trade system because it lets the market set the emissions price. Setting the right emissions price, or cap level, is complicated for the climate change issue by the presence of deep uncertainty about the marginal damage of emissions as well as whether and at what emissions level potential catastrophes are possible.

In attempts to get the best of both tax and cap approaches, or at least reach political compromise, hybrid policies have been proposed—a cap-and-trade system with a "safety valve". Under this hybrid system, if the carbon price rises above a predetermined level (the safety valve price) then the government will sell an unlimited amount of emissions allowances at that safety valve price. If the safety valve is triggered, the emissions cap is no longer met and the policy effectively becomes a carbon tax at the level of the safety valve price. Whether the safety valve is likely to be triggered is influenced by the level of the cap and by the safety valve price. The more stringent the cap, the higher the expected carbon price in the absence of a safety valve. If the safety valve price is set relatively high in relation to the expected carbon price, then resorting to government sales would be less likely. If the safety valve is set relatively low in relation to the

142

expected carbon price, then use of the safety valve is more likely. A tight cap with a relatively low safety valve will mean the cap very likely will be exceeded. If the safety valve price reflects the social cost of carbon, the policy would achieve the same efficiency as a carbon tax.

A safety valve has been included in several proposed U.S. cap-and-trade legislative bills, under different names. The Bingaman-Specter bill (S.1766) of 2007 incorporated a safety valve, called the Technology Accelerator Price (TAP), which allowed firms to pay the TAP in lieu of submitting permits for their emissions. The Waxman-Markey and Kerry Boxer bills included a Strategic Allowance Reserve that allowed the U.S. EPA to sell permits into the carbon market via an auction with a set minimum price in the event permit prices rose faster than expected. The Lieberman-Warner bill included a Carbon Market Efficiency Board that would monitor the permit market and release permits when the cost got too high. The Cantwell-Collins bill included a "price collar"—a price floor and a price ceiling—on auctioned permits.

Because this hybrid policy approach of a cap-and-trade system with a safety valve has been frequently proposed, it is useful to know what effect it has on near-term low-carbon generation and emissions reduction decisions. Theoretical and modeling studies suggest that a safety valve would provide flexibility that can protect against unexpectedly high costs or price shocks of a cap-and-trade system (e.g. Weitzman, 1974; Roberts & Spence, 1976; Pizer, 2002; Pizer, 2005; Jacoby & Ellerman, 2004; Webster *et al.*, 2008, Ellerman *et al.*, 2008). However, safety valves have also been criticized, particularly by environmental groups, for allowing emissions to rise above the cap. Modeling studies have also predicted that safety valves could discourage investment in new technologies to reduce emissions (e.g. Burtraw & Palmer, 2006).

The hypothesis for this work is that the existence of a safety valve will undermine incentives to make low-carbon investments in the near-term and thus make it more likely to trigger the safety valve in the longer term. Thus, unless the safety valve is so high that there is little chance of it ever being triggered, incorporating one in a policy makes it almost certain that it will be triggered. The results below investigate this hypothesis.

7.1.2 Modeling a Safety Valve

In this section, the policies limiting emissions include a safety valve. The three possible policies are: (1) no policy, (2) an emissions cap of 20% below no policy emissions with a safety valve, and (3) an emissions cap of 40% below no policy emissions with a safety valve. The

safety valve gives the option of paying a fixed price per ton of CO_2 (the safety valve price) for each unit of emissions *over* the cumulative cap rather than abating emissions to be under the cap. This chapter explores different safety valve prices, starting with \$52/tonCO_2, which is the latest U.S. government estimate of the 2030 social cost of carbon (U.S. Government, 2013).[18]It is assumed that the safety valve price is known in Stage 1. For example, in Stage 1 we know that one of the possible policies is a -40% cap with a \$52 safety valve and another possible policy is a -20% cap with a \$52 safety valve. The three possible low-carbon technology costs are also considered: the Stage 2 cost of low-carbon generation relative to conventional generation (i.e. the cost markup) will either be 3, 1.5, or 1. The three cost markups are assigned the Pearson-Tukey probabilities of 0.185, 0.63, and 0.185, respectively (see Section 6.1). In this section there is no technological learning for the low-carbon technology cost. The uncertainty in the low-carbon technology cost contributes to uncertainty in the carbon price and policy cost, and in turn whether or not the safety valve will be triggered.

The safety valve is implemented within the CGE model. If the carbon price rises above the safety valve price, additional emissions permits are made available so that the resulting carbon price equals the safety valve price. In this way, emissions over the original cap are taxed at the safety valve price.

Based on the emission limit and the carbon price that results given the Stage 1/Stage 2 strategy, the DP then determines whether it is worth it to pay the safety valve penalty to exceed the cap or to pay the additional abatement cost that would be needed to reduce emissions enough to meet the cap. Through backward induction, the model identifies the best Stage 1 decisions. We can then explore how the first stage decision may change when there is a safety valve in the second stage policy, allowing the possibility of exceeding the cap, but at a cost.

7.2 Results

First we investigate scenarios in which the policy is known for certain to be a -40% cap, but the low-carbon technology cost is uncertain. We compare the case when the cap is strictly enforced with policies that include a safety valve at different price levels. Figure 7.2 shows the impact of a safety valve on the optimal Stage 1 electricity investment mix under a deterministic

---

[18] Estimate is in 2007 USD and used a discount rate of 3%.

policy. A safety valve price of \$52/ton CO_2 results in no low-carbon investment in Stage 1. Even a \$200/ton CO_2 safety valve results in no near-term low-carbon investment, though it does encourage some shifting from coal to natural gas generation. It takes a safety valve price of approximately \$300/ton CO_2 to encourage the same electricity investment mix chosen when the cap is enforced with certainty. The safety valve price needed to spur significant low-carbon investment in Stage 1 needs to approach the expected 2030 carbon price that would result if no low-carbon investment is made in Stage 1 but the cap is enforced
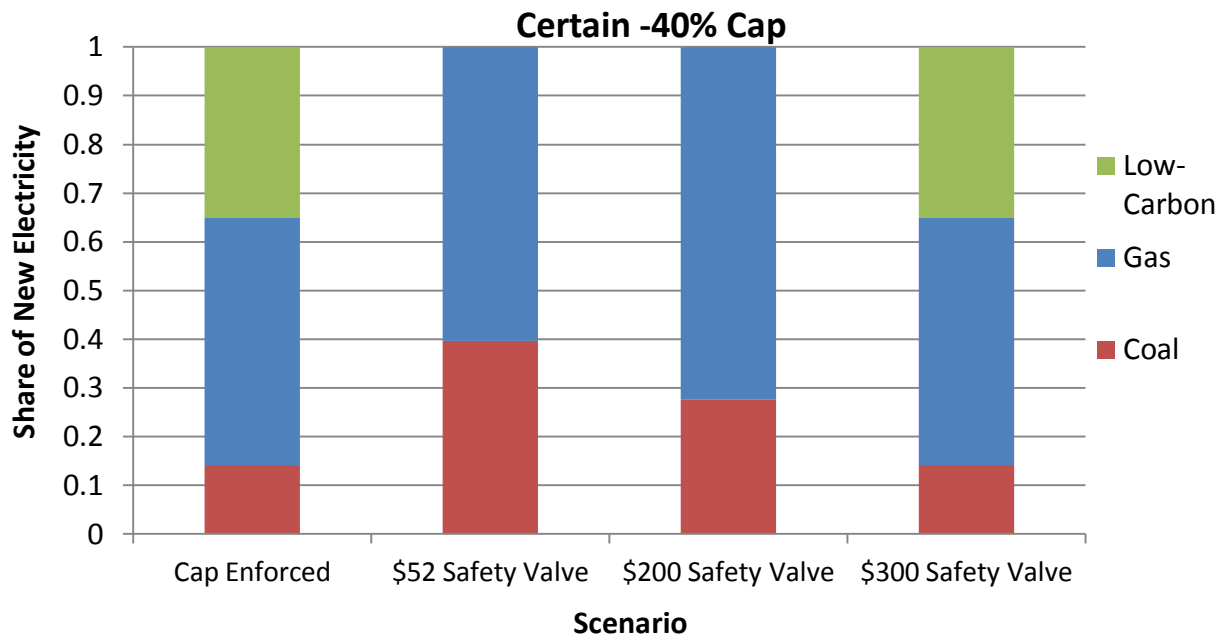
Figure 7.2 Stage 1 Shares of New Electricity under Deterministic -40% Cap with Safety Valve

Table 7.1 shows expected 2030 carbon prices that result from different policy and technology cost scenarios when the cap is enforced (i.e. there is no safety valve) and either the optimal hedging strategy (as identified in the previous chapter) is pursued in Stage 1 or no low-carbon investment is made in Stage 1. If no policy is implemented, there of course is no carbon price. The carbon prices that result when no near-term investment is made are much higher than those that result when pursuing the optimal hedging strategy. In order to encourage near-term action, the safety valve price needs to approach the carbon price that would result from making no low-carbon investment in Stage 1 and then having to meet the cap. Otherwise, it is better to do nothing in Stage 1, leading to high Stage 2 carbon prices which trigger the safety valve, thereby allowing emissions to exceed the nominal cap (see Figure 7.3). Only if the safety valve

price is high enough to encourage enough Stage 2 reductions to meet the cap anyway, does it make sense to start investing in Stage 1 and replicate the optimal hedging strategy. Lower safety valve prices discourage near-term action because they will be triggered, allowing emissions to exceed the cap. So the safety valve contains policy costs by requiring fewer emissions reductions. It is better (i.e. results in lower policy costs) to (a) incur no extra investment costs in Stage 1, pay the high safety valve price, and reduce far fewer emissions than the cap requires, than to (b) incur extra Stage 1 investment costs, pay the lower carbon prices, but reduce enough emissions to meet the cap. Ultimately, fewer reductions at a high safety valve price leads to less total consumption loss than more reductions at a lower carbon price.

Table 7.1 Expected 2030 Carbon Prices when the Cap is Enforced (No Safety Valve): Optimal Stage 1 Strategy vs. No Stage 1 Low-Carbon Investments

	Expected 2030 Carbon Price ($/tonCO2)	
	Cap Enforced: Optimal Stage 1 Strategy	**Cap Enforced: No Stage 1 Low-Carbon Investments**
Certain -40% Cap		
MU3	413.12	459.21
MU1.5	31.61	337.51
MU1	0.00	300.83
Expected	96.34	353.24
Certain -20% Cap		
MU3	191.47	305.27
MU1.5	56.63	169.16
MU1	0.00	84.37
Expected	71.10	178.65
Policy Uncertainty (1/3 Probability of Each Policy)		
Then -40% Cap, MU3	457.87	459.21
Then -40% Cap, MU1.5	129.91	337.51
Then -40% Cap, MU1	0.00	300.83
Then -20% Cap, MU3	117.65	305.27
Then -20% Cap, MU1.5	29.43	169.16
Then -20% Cap, MU1	0.00	84.37
Expected	67.72	177.30

Figure 7.3 Cumulative Emissions under Deterministic -40% Cap with Safety Valve

For the -40% cap, when enforced, the expected carbon price under the optimal hedging strategy is under $100, while it is about $350 if no near-term investment is made. Accordingly, the safety valve price needs to be in the $300 range to encourage the same amount of near-term low-carbon investment as when the cap is enforced with certainty. At lower safety valve prices, no near-term investment is optimal as it guarantees that the safety valve will be triggered, reducing the effective stringency of the policy. Even though the $200 safety valve price is higher than the expected carbon price that results from following the optimal hedging strategy it is worth it to pay the high safety valve price in order to emit above the cap. Making no near-term investments guarantees that safety valves below about $300 will be triggered, thereby effectively reducing the stringency of the policy.

Next we investigate scenarios in which the policy is uncertain with each policy being equally probable and the low-carbon technology cost is also uncertain (Figure 7.4). The resulting 2030 carbon prices depend on which policy is implemented and which technology cost is realized (Table 7.1). Notice in Table 7.1 that under policy uncertainty, if the policy turns out to be the -40% cap, the 2030 carbon price is higher than the carbon price when the same -40% cap is known with certainty (for example, $130/tonCO_2$ vs. $32/tonCO_2$ when the markup is 1.5).

147

This is due to the difference in Stage 1 strategies with and without policy uncertainty. As we saw in Chapter 5, when it is known with certainty there will be a -40% cap, more low-carbon investment is optimal in Stage 1 compared to when the policy is uncertain. For a given emission reduction target, more low-carbon generation leads to lower carbon prices in Stage 2: emissions reductions from low-carbon generation means fewer reductions must occur from conventional generation, so there is less demand for carbon permits and therefore a lower carbon price (see, for example, Morris et al., 2010). Generally, the less near-term action taken, the higher the Stage 2 carbon price will be to meet the policy.

Figure 7.4 Share of New Electricity under Policy Uncertainty with Safety Valve

The first case considered in Figure 7.4 is when we do not know what the policy will be, but we know it will be enforced with certainty. In that case, 20% low-carbon investment is optimal. The remaining cases in Figure 7.4 are ones in which we do not know what the policy will be, but we do know it will include a safety valve at a known price (e.g. $52, $200 or $300/ton CO_2). As with the deterministic policy, we see that the inclusion of a safety valve reduces the optimal Stage 1 low-carbon investment. Under policy uncertainty it takes a safety valve price of $300/ton CO_2 for investors to choose the same investment mix chosen when there is no safety valve. Even though the $200 safety valve price is higher than the expected 2030 carbon price that results when taking no near-term action ($177), there is still a significant

148

chance (about 40%) of a $200 safety valve being triggered (i.e. several policy/technology cost outcomes result in carbon prices above $200). There is only a 6% chance of carbon prices rising significantly above $300, and as a result a $300 safety valve is high enough to encourage the same amount of near-term action as when the cap is enforced. Again, at high enough safety valve prices that have low chances of being triggered, it is better to pursue the optimal hedging strategies to meet the cap. However, at lower safety valve prices, it is better to make no near-term low-carbon investments and instead count on the safety valve being triggered in order to be able to emit over the cap. As Figures 7.5 and 7.6 show, this strategy results in cumulative emissions above the caps. If the policy turns out to be a -40% cap, emissions greatly exceed the cap—by 50% with a $52 safety valve and by 33% with a $200 safety valve. If the policy turns out to be a -20% cap, emissions still exceed that cap—by 12% with a $52 safety valve and by 2% with a $200 safety valve.

Figure 7.5 Cumulative Emissions when Uncertain Policy Then -40% Cap with Safety Valve

Figure 7.6 Cumulative Emissions when Uncertain Policy Then -20% Cap with Safety Valve

In general, safety valve prices that are lower than the carbon prices that would result if no low-carbon investment is made in the near-term when the cap is strictly enforced have the effect of reducing the stringency of the policy. Under uncertainty in policy with low safety valves, investors are essentially facing three possible policies that are less stringent than the original policies (without safety valves). In the face of less stringent future policies, investors would choose to do less (less low-carbon generation and fewer emissions reductions) in both Stage 1 and Stage 2 than they would if they knew the caps would be enforced. The safety valve must be high enough so there is little chance of it ever being triggered, even if no near-term action is taken, in order to incentivize the same strategy as when there is no safety valve. Any lower safety valve price will discourage low-carbon generation and emissions reductions investments in Stage 1 and Stage 2.

Unless the safety valve price is sufficiently high that it is unlikely to be triggered even if no near-term low-carbon investment is made, including a safety valve in the policy undermines the ability to meet the cap. There is a moral hazard: knowing that the cap can be loosened if the policy becomes too costly acts as insurance against high policy costs and discourages investment in new technologies to reduce emissions and avoid high policy costs. Investors lose incentive to

act now since the cap may not even be enforced in the future. Less aggressive action in the near-term makes it more difficult to meet the cap in the future, leading to high carbon prices which cause the safety valve to be triggered and emissions to exceed the cap. In this way, the safety valve is a self-fulfilling prophecy: including a safety valve to guard against high policy costs leads to investment decisions that actually cause high policy costs and necessitate the use of the safety valve, and in turn the violation of the cap.

As an example, let us consider the $52 safety valve case when the policy and technology cost is uncertain (as in Figure 7.4). If investors pursued the Stage 1 investment strategy from the $52 safety valve case (i.e. no near-term low-carbon investment), how costly would it be to actually meet the cap that is ultimately implemented? How does that compare to the cost of meeting the cap when following the Stage 1 strategy chosen when there is no safety valve? If the three policies are equally likely, the optimal Stage 1 strategy without a safety valve is 20% low-carbon, 17% coal and 63% natural gas generation, and 18% emissions reductions. If there is a $52 safety valve price, the optimal Stage 1 strategy is 40% coal and 60% natural gas generation and no emissions reductions. Table 7.2 provides the policy cost comparisons from following these two strategies. If the cap were actually enforced in all cases, the policy costs (reduced economic consumption relative to a reference no policy case) would be much higher if the strategy from the $52 safety valve case was pursued in Stage 1 than if the optimal strategy without a safety valve was pursued. If the cap turns out to be 40%, following the $52 safety valve strategy increases the cost of meeting the policy by 1023% ($2322 billion). If the cap turns out to be a 20% cap, the $52 safety valve strategy increases policy costs by 159% ($166 billion). If there turns out to be no cap, the $52 safety valve strategy results in policy costs that are 100% ($72 billion) lower. Overall, following the $52 safety valve strategy increases the expected cost of meeting the future cap by 601% or $806 billion. Of course, this increased cost is not truly realized because the safety valve is triggered and emissions rise above the cap that is implemented. However, it is noteworthy how the inclusion of a safety valve leads to different investment decisions that make it more costly to meet the cap, hence the self-fulfilling prophecy.

Table 7.2 Policy Costs of $52 Safety Valve (SV $52) Strategy vs. No Safety Valve (No SV) Strategy if Cap Enforced in All Cases

Scenario: If cap were enforced in all cases	% Change Consumption from Reference	Change Consumption from Reference in $billion	% Change Policy Costs SV $52 vs. No SV	Change Policy Costs SV $52 vs. No SV in $billion
SV $52 then -40% Cap	-2.71%	-2549	-1023%	-2322
No SV then -40% Cap	-0.25%	-227		
SV $52 then -20% Cap	-0.29%	-270	-159%	-166
No SV then -20% Cap	-0.11%	-104		
SV $52 then No Policy	0.00%	0	100%	72
No SV then No Policy	-0.08%	-72		
SV $52 Expected	-1.00%	-940	-601%	-806
No SV Expected	-0.15%	-134		

It is useful to help unpack why the different Stage 1 strategies result in such different policy costs if the cap is enforced. The pie charts in Figure 7.7 show how the cumulative emissions required for the -40% cap are divided between Stage 1 and 2 for (a) the no safety valve case and (b) the $52 safety valve case. When there is no safety valve, the optimal Stage 1 strategy includes more low-carbon generation and more emissions reductions, so Stage 1 takes up less of the total allowed emissions and more emissions are available in Stage 2. When there is a $52 safety valve, the optimal Stage 1 strategy includes no low-carbon generation or emissions reductions, so Stage 1 takes up more of the total allowed emissions, leaving fewer for Stage 2. With fewer emissions available in Stage 2, if the cap is to be met, the carbon price will be higher (see Table 7.1), as will the policy cost. These higher carbon prices in the $52 safety valve case trigger the safety valve, causing emissions to rise above the cap.

Thus, inclusion of the safety valve causes less aggressive action in the near-term, making it more difficult to meet the cap in the future, leading to high carbon prices which cause the safety valve to be triggered and emissions to exceed the cap. Another way to think about this is that actions today can affect the stringency of the policy in the future. With a safety valve, doing less now results in future use of the safety valve, which in effect loosens the cap. Doing more now, reduces the cost of meeting the policy, removing or reducing the need to use a safety valve and allowing a more stringent cap to be met.

(a) No Safety Valve

- ■ Stage 1 Cumulative Emissions
- ■ Stage 2 Cumulative Emissions

(b) $52 Safety Valve

- ■ Stage 1 Cumulative Emissions
- ■ Stage 2 Cumulative Emissions

Figure 7.7 Cumulative Emissions in Stage 1 and 2 for -40% Cap (a) When No Safety Valve and (b) When Safety Valve of $52/tonCO_2$

7.3 Discussion

This chapter has important implications for policymaking. First, the safety valve represents a tradeoff between policy costs and emissions (and the costs associated with emissions). This tradeoff is represented in Figure 7.8. The higher the safety valve price, the lower the emissions, but the higher the policy costs. In turn, lower safety valve prices result in lower policy costs but more emissions. Policymakers must think carefully about this tradeoff.

What are their goals in terms of emissions and what cost are they willing to accept to meet those goals?

Figure 7.8 Emissions vs. Policy Costs: Uncertain Policy with Safety Valve

Second, policymakers must take great care when designing policy since the inclusion of a safety valve can undermine the policy. If the goal is to truly meet an emissions cap, a safety valve may not be desirable or, if included, the safety valve price must be carefully set. It should be set high enough that it is truly a "safety valve"- a way to avoid excessively high unexpected costs of meeting the policy. Setting the safety valve price too low would simply reduce the de facto stringency of the cap. Careful analysis should be conducted and the safety valve price should be set above the expected carbon price for the policy. If policymakers give the impression that they are serious about enforcing any future cap, investors will be encouraged to take more aggressive actions today, which will ultimately help ensure that the future cap can in fact be met without costs being much higher than expected.

Another consideration raised by these results is that of an *implicit* safety valve in the form of the government loosening or canceling the policy in face of high costs. The above results show that explicit inclusion of a safety valve leads to less near-term investment, increasing the probability that the safety valve will be triggered and the cap will be exceeded. One might expect

154

this type of behavior to also hold to some extent in the presence of an implicit safety valve: investors know that the government will not enforce a policy regardless of the cost, and therefore pursue less aggressive action with the expectation of being "bailed out". Such an environment is amenable to game theory analysis. Firms and the government could be seen as essentially playing "chicken" over a policy—will firms give in and do more to try to meet the policy or will government give in fearing high policy costs and loosen/cancel the policy? This could be investigated in future work.

In the next and final chapter, key insights, implications and contribution of this dissertation are summarized. Relevant limitations and opportunities for future research are also discussed.

(This page is intentionally left blank)

Chapter 8: Conclusion

This chapter presents final remarks on the dissertation and its findings. Section 8.1 provides a brief summary of this research and the tasks undertaken. Section 8.2 offers key insights from the analyses conducted. Section 8.3 discusses the implications of the results for policy. Section 8.4 reviews technical and modeling contributions of this work. Section 8.5 reflects upon many of the limitations of this work, outlining future research opportunities, and Section 8.6 offers final thoughts.

8.1 Summary of Research

In order to provide insight to near-term electricity investment strategies under uncertainty, this research undertakes several tasks. A computable general equilibrium (CGE) model focused on electricity generation is developed. A stochastic dynamic programming (DP) component is then created to link with the CGE model to study the impact of uncertainty. The uncertainties of interest are uncertainty in future policy limiting CO_2 emissions and uncertainty in the future cost of low-carbon generation technologies. A small survey is conducted to gain understanding about expectations about policy that exist in the electricity industry. Literature on technology cost uncertainty is reviewed to inform distributions of the future cost of low-carbon technologies. Also, a new stochastic formulation of technological learning is developed, informed by the learning-by-doing literature. The new DP-CGE model is applied to investigate a number of factors affecting optimal near-term electricity investments: (1) policy uncertainty, (2) expansion rate limits on low-carbon technologies, (3) low-carbon technology cost uncertainty, (4) technological learning, and (5) the inclusion of a safety valve in future policy. The results of this research provide several key insights and policy implications, and the modeling framework provides technical and modeling contributions.

8.2 Key Insights

The optimal level of near-term investment depends on which uncertainties one considers and the magnitude and distributions of those uncertainties. Uncertainty creates risks, and the best one can do in the face of uncertainty is to develop near-term investment strategies that hedge against those risks. In this work, the two risks considered are: (1) the sunk cost of near-term low-

carbon technology investments that prove to be unnecessary and (2) facing very high future policy costs. The first risk is created if near-term investments in expensive low-carbon technologies are made but ultimately no policy limiting electricity emissions is implemented, or if a weak policy is implemented that does not require low-carbon technologies to the extent they were deployed. Similarly, even if there is a future policy, near-term investments could be wasted if they do not decrease the cost of the technology in the future and that the technology remains expensive. The second risk is created when future policy requires significant emissions reductions and significant amounts of low-carbon generation, but failure to make (sufficient) near-term investments results in high costs to meet the policy. A carbon-intensive electricity mix combined with expensive low-carbon technologies that may not be able to expand as much as desired leads to high policy costs in that case. How to balance these two risks depends on the perceived probabilities of future outcomes and requires explicitly considering the uncertainty in the decision-making process. This work shows that there is value to investing in low-carbon generation now and taking on some of the first risk in order to provide flexibility in meeting future policy and reduce the second risk.

The results from this work suggest that in the presence of uncertain future policy and technology costs there are three key motivations for near-term investments in low-carbon technologies beyond what is needed to meet near-term goals alone.

First, near-term investments lower the expected cost of emissions reductions in the future. Current investments create a less carbon-intensive electricity mix, which would make meeting a future emissions reduction policy easier. Spreading the burden of emissions reductions over time by pursuing some reductions now in anticipation of a future limit on emissions reduces the expected discounted cost of meeting a policy.

Second, near-term investments in low-carbon technologies may lower the expected cost of those technologies in the future, through technological learning and economies of scale effects, which in turn lowers the expected cost of emissions reductions. Large shares of low-carbon generation may be required in the future, depending on the stringency of future policy, and investments now to make those technologies cheaper would make meeting the policy less costly. Having low-cost low-carbon technologies available in the future provides flexibility in how to achieve future emissions reductions if needed.

158

Third, if there are constraints on the expansion rate of low-carbon generation, creating the needed infrastructure for commercialization of these new technologies now can help limit bottlenecks to greater expansion in the future if it turns out that the policy is stringent enough to require significant use of low-carbon technologies. Factors that may limit the expansion rate of a technology include limited trained engineering capacity and electricity system adjustment costs. Near-term investments could help overcome those limits, providing flexibility to rapidly expand the future use of low-carbon technologies if needed in order to meet a future emission limit.

Overall, more near-term low-carbon investments create future flexibility, which justifies the additional near-term cost of investments. The value of this flexibility is only explicitly considered in the context of decision-making under uncertainty.

There are three main approaches to handling uncertainty. First, we can ignore uncertainty, as is done in the two prevailing frameworks for economic models. Economic models with myopic expectations assume nothing will change in the future, while models with perfect foresight assume we know the future with certainty. Both of these ignore uncertainty. Second, we can acknowledge uncertainty and assess different possible scenarios as if they are certain (using either of the economic frameworks above). We can then take a "middle of the road" or "average" approach and use the middle scenario as certain or use the average of the uncertain value as certain. Monte Carlo analysis falls into this second category because each draw of a parameter value from a distribution is used in the underlying model as if it is known to be the true value. A third option is to formally represent decision making *under* uncertainty (by using a dynamic programming framework, for example), which makes use of the imperfect expectations we have about the future to develop hedging strategies and allows for learning and revising decisions over time. This research shows the value of this third approach.

This dissertation demonstrates that formally considering uncertainty (and the risks it creates) in the decision-making process results in different near-term decisions than would have been made if the uncertainty were ignored or not truly incorporated into decision-making. Considering decisions under uncertainty results in a hedging strategy: invest in more low-carbon generation now than would be needed for near-term goals alone in order to minimize expected future policy costs. The amount of near-term investment depends on the expectations of future outcomes and on the expected impact of near-term investments on future costs. More near-term investment should be made as we expect stringent policy to be more likely. As the expected

probability of a stringent cap increases, the potential need for significant amounts of low-carbon generation to meet the cap increases, and in turn the value of near-term investments in those technologies increases because they reduce the expected discounted cost of meeting the policy. They do so by reducing the expected future cost of low-carbon technologies, helping to overcome technology expansion rate constraints, spreading the burden of emissions reductions over time and developing a less carbon-intensive electricity mix—all of which provide future flexibility. In addition, as long as there is a chance of stringent policy that will necessitate the use of low-carbon technologies, the amount of optimal near-term investment increases with the magnitude of the learning effect on the technology costs as well as with constraints on the rate of technology expansion. If technological learning from near-term investments is more effective at decreasing the expected future cost of low-carbon technologies, the value of near-term investments increases because you receive "more bang for your buck". Similarly, the stricter the constraints on the rate of technology expansion, the greater the value of near-term investments. Near-term investments can help overcome expansion rate constraints and provide flexibility to use as much of the technology in the future as desired.

Although the optimal amount of near-term investment in low-carbon technologies changes with the perceived probability of outcomes, general insights can be drawn. Across probability distributions, it is clear that near-term investments in low-carbon technologies should be greater than what would be considered cost-effective in the near-term alone (i.e. without considering how uncertain future outcomes may impact current investments). Also, overinvestment in conventional technologies poses a greater risk under policy uncertainty than overinvestment in low-carbon technologies due to differences in variable costs. If a stringent policy is implemented, overbuilt conventional capacity may need to be stranded due to high operating costs driven by high fuel costs that reflect the carbon price. If no or weak policy is implemented, overbuilt low-carbon capacity means unnecessary sunk capital costs, but low operating costs allow the capacity to continue to be used.

In addition, a review of the literature revealed typical learning rates for low-carbon technologies to be 20-25%, but this work demonstrated that even lower learning rates of 10-15% can lead to significant additional near-term investment in the low-carbon technology in order to lower the expected future cost of the technology in case a stringent policy is adopted.

Further, if there is a threshold of low-carbon technology capacity beyond which there are no longer adjustment costs that limit the rate of expansion of the technology, reaching that threshold in the near-term can minimize expected policy costs by providing maximum flexibility for future expansion of that technology as needed in the future. To that end, identifying such thresholds would be a useful research initiative that could inform near-term investment decision.

This work also shows that uncertainty has a cost, beyond the cost of meeting the policy. In the experimental design used here, uncertainty in the future policy increases the expected discounted cost of policy by over 45%. Even while pursuing the optimal hedging strategy under uncertainty, once the policy is known, those hedge decisions are not necessarily optimal in retrospect. Had the policy been known for certain, different near-term decisions would have been made that would have resulted in lower policy costs. We can think of this as the cost of political indecision and/or of missing scientific information, which are the two main drivers of policy uncertainty. While regularly changing and divergent politics are partly responsible for policy uncertainty, so too is the uncertainty in the science of climate change, which makes it difficult to determine the level of emissions reduction action that should be taken. While these drivers are beyond the scope of this research, the high cost of uncertainty suggests the value of doing what we can to resolve it—in particular, more climate research (to the extent that it can improve our understanding of the severity of the problem) and efforts to develop a political consensus sooner. If consensus can be reached and the science uncertainties resolved, setting clear, long-term policies can minimize expected policy costs. Given the risks associated with climate change as reported in various national and international research assessments, such as the Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC, 2007) and the U.S. Second National Climate Assessment (USGCRP, 2009), it would seem difficult to make the case that no emissions reduction policy is warranted now or in the future.

8.3 Policy Implications

Before considering the implications of the results for policy, we must first place the model used in the context of the real world. As with any model, simplifications are made that may affect how the results are translated to a real world policy setting.

This work is focused on societal costs and identifying strategies that maximize social welfare. In doing so, we assume that private entities, like electric utilities, act in accordance with

the goal of maximizing social welfare. A market system can deliver socially optimal results under some circumstances—perfect competition and no externalities. In other circumstances, government policy may be needed to achieve the socially optimal strategy. In this work, future climate policy would internalize the negative externality of carbon emissions, which should cause private actors to look forward and act on that expectation in a way that is consistent with the socially optimal strategy. However, an additional externality could deter such behavior by private actors—the positive externality caused by technological learning spillovers. Learning spillovers create a positive externality in which actors other than those making the investment benefit from it without sharing in the cost. This results in a market failure in which the positive externalities are not factored into private investment decisions and the result is an underinvestment in low-carbon technologies relative to what is socially optimal. In the model, the benefits of technological learning are fully captured and taken into account by the central planner in identifying the optimal near-term strategy. However, in a more realistic industry structure where there is competition and potentially technological learning spillovers, there may be need for government policy to encourage private investment in low-carbon technologies.

Accordingly, in a world with technological learning spillovers, government policy may be needed to force investment closer to the socially optimal hedging strategy. Near-term policies encouraging investments in low-carbon technologies prior to the implementation of an emissions cap could help reduce the expected cost of meeting the future cap. Examples of such policies include a renewable or clean portfolio standard, which requires that a certain share of electricity generation comes from renewable or other clean technologies, tax incentives for clean technologies, and demonstration projects of clean technologies. Such policies would encourage or require investors to pursue a hedging strategy, which could help lower the expected cost of low-carbon technologies in the future and could provide greater flexibility and ease in meeting a future emissions cap. This of course assumes that policymakers have a better understanding of the probabilities of future emissions limits than investors do and can identify the socially optimal hedging strategy and then implement a policy to achieve that strategy. If this assumption is not true, government policy of this form could increase costs for private actors, for example by requiring more investment in low-carbon technologies than is necessary to minimize expected policy costs. Nevertheless, it remains a robust finding of this work that if future emissions

restrictions are possible, some level of near-term development of low carbon technologies is socially beneficial.

Another policy implication from this work is that the inclusion of a safety valve within an emissions cap policy should be considered with care as doing so can undermine the ability to meet the cap. The existence of the safety valve makes it more likely that the cap will not be met due to a moral hazard: the safety valve acts as insurance against high policy costs and therefore discourages near-term low-carbon investments to try to avoid high policy costs. Specifically, unless the safety valve is so high that there is little chance of it ever being triggered, even if no near-term action is taken, the inclusion of a safety valve leads to less aggressive action in the near-term, making it more difficult to meet a cap in the future, leading to high carbon prices and compliance costs, which cause the safety valve be triggered and emissions to exceed the cap. In this way, the safety valve is a self-fulfilling prophecy: including a safety valve to guard against high policy costs leads to investment decisions that actually cause high policy costs and necessitate the use of the safety valve. Setting an appropriate safety valve price is difficult, and subject to political debate. Theoretically it should be set to the social (damage) cost of carbon (however, if we knew the social cost of carbon, it would be most efficient to use a carbon tax at that level to reduce emissions). Regardless, the result of this research is that if we set a modest safety valve price, then forward-looking behavior will almost assure that it will be triggered.

8.4 Technical and Modeling Contributions

The main technical contribution of this dissertation is to provide a new decision support framework that considers economy-wide effects of electric power sector investment decisions under uncertainty. There have been sector-specific studies that capture decision-making under uncertainty well, but cannot address economy-wide social welfare implications. There have been economy-wide computable general equilibrium (CGE) studies with uncertainty, but without capturing the critical nature making decisions under uncertainty, learning, and then making decisions again. This work makes the unique contribution of modeling decision-making under uncertainty with learning and the ability to revise decisions over time in a framework that represents the entire economy and can measure social welfare impacts. Applying this new modeling approach allows for a quantitative evaluation of near-term electric power investment decisions under uncertainty. In doing so, this work demonstrates how a CGE model can be

163

restructured to capture stochastic dynamic expectations. As a result, this work facilitates further model development and analysis in this area. In addition, the modeling approach developed here can be applied to other sectors, other decisions, and other uncertainties.

This dissertation also demonstrates the value of formally representing uncertainty in decision support models. As noted previously, existing modeling approaches typically do one of two things with regard to uncertainty: (1) ignore it or (2) consider scenarios, each with certainty, and then use the middle one or use the average of the uncertain value as though certain. These approaches result in different investment strategies than the optimal hedging strategy from a dynamic programming framework that explicitly considers decision making under uncertainty. In the experimental design used in this work, the near-term investment strategy from the average scenario increases the expected policy cost by over 50% compared to the optimal hedging strategy. Investment strategies identified when ignoring uncertainty also increase the cost of meeting policy. Most economic models fall into one of two extremes in terms of expectations, both of which ignore uncertainty. One is myopic expectations, which assumes nothing will change in the future. Near-term strategies developed under this framework will not make any investments that are not cost-effective in the near term. As a result, such a strategy increases the expected policy costs by more than four times the cost of the optimal hedging strategy. The other type of expectations is perfect foresight, which assumes we know with certainty how the future will unfold. If we are wrong about the future, the strategy chosen may be costly. The investment strategy developed under this framework and how the resulting policy costs compare to the optimal hedge strategy depends on what future is assumed and how that compares to the probabilistic future in the dynamic programming framework. In the experimental design from this work, assuming perfect foresight increases the policy cost by anywhere from 12% to over 400% depending on the future emissions limit assumed. Perfect foresight strategies that assume a limit on emissions fare better than the myopic strategy because it is more costly to overinvest in conventional generation than to overinvest in low-carbon generation in the face of policy uncertainty. Overall, this work demonstrates that considering uncertainty in decision-making results in investment strategies that best minimize expected policy costs.

In addition, this dissertation exemplifies the value of an engineering systems research approach (de Weck *et al.*, 2011). This work pursued a system-level analysis capturing interactions of the electric power sector, the overall economy, and policy. It did so by integrating

three distinct disciplines: electric power planning, economic modeling, and sequential decision analysis. Such an approach allows the study of a complex, large-scale problem in a new way that provides new insight.

Further, this work presents a stochastic learning-by-doing approach, contributing to the literature in that area. The unique stochastic formulation of technological learning has near-term investment in a technology affect the probability distribution of the future cost of that technology. This formulation captures the fact that the impact of investments on technology cost is uncertain.

8.5 Limitations and Future Work

In developing the model used for this work, the goal was to keep the model as simple as possible while also capturing key dynamics. A small, transparent model is particularly valuable for the purposes of demonstrating the role of uncertainty in decision-making. However, this approach necessarily means that several details that could make the model more useful as an applied decision support tool have been purposefully omitted. This section reviews some of the limitations of the current model and discusses opportunities for future research.

Although the CGE component of the model captures complex dynamics, it is highly aggregated. The representation of a single conventional generation technology and a single low-carbon technology was suitable for providing insight on the impact of uncertainty on investments, but fails to provide the level of detail required by investors. In particular, representing a separate natural gas technology and multiple low-carbon generation options (e.g., wind, solar, nuclear, CCS, etc.) would more fully reflect the investment choices available. The addition of more technology options would increase the size of the dynamic program by increasing the decision space. Instead of choosing the share of low-carbon generation and the model finding the complementary share of conventional generation, decisions would be explored for the share of all but one of the technologies. This increased technological detail would allow for exploration of optimal generation portfolio decisions. In addition, the vast majority of the economy is aggregated into the "other" sector. Disaggregating into transportation, energy-intensive industries, service, agriculture, and other sectors would more accurately reflect the composition and interactions of the economy. Doing so would also allow investigation into other sectors under uncertainty. Increasing sectoral disaggregation would involve adjusting the social

165

accounting matrix. Overall, an increased level of detail would improve the usefulness of the model as a decision support tool.

In addition to the high level of technology and sectoral aggregation, the CGE model represents a single representative agent with a set of expectations about the future. Of course in the real world, there are many distinct actors, all operating under expectations about the future outcomes of uncertainty. The survey described in Chapter 5 in fact exemplifies very different expectations about future policy. Because optimal near-term strategy is driven by the expectations about the future held by those making the investments, investors with different expectations would make different near-term decisions. Representing multiple agents would allow exploration of how the aggregate of diverse private investment decisions diverges from the socially optimal investment strategy. One could also study how policy designed to achieve the socially optimal investment strategy impacts the different actors represented.

Different objective functions in the dynamic program could also be explored. In this research the objective is to choose actions to maximize total expected discounted social welfare in the economy over the planning horizon. Under policy uncertainty, this is equivalent to minimizing expected policy costs. However, other objective functions could be explored. The objective could be to reduce the probability of right-hand tail (i.e. high cost) outcomes or to minimize the maximum possible policy cost. Different objective functions would likely lead to different optimal strategies.

For the dynamic programming component of the model, the decision and uncertainty space is discretized at a level suitable for the purposes of demonstration and insight, but potentially unsatisfactory for real world decision making. The decision space—low-carbon generation share and emissions reduction level—is modeled at fairly high resolution—both in steps of 5%. The uncertainty space is modeled at a low resolution—discrete three-point probability distributions are employed. Fuller probability distributions with more potential future outcomes would be more realistic. Similarly, it would be more realistic if the two-stage decision period framework was extended to multiple decision periods, perhaps in five-year time steps, and covering a longer time frame. This would demonstrate the continual learning and updating of expectations over time. Of course the resolution of each of these areas could be increased, but at the cost of computational complexity and time.

Future research could apply the framework developed here to a more detail-rich CGE model and a higher resolution dynamic program. In particular, utilizing the MIT EPPA model for the CGE component would allow detailed analysis of energy policies and decisions under uncertainty. Significantly scaling up the CGE model and/or the dynamic programming resolution would cause the combined model to run into the "Curse of Dimensionality" in which it would become computationally prohibitive to solve the full dynamic program. Instead, approximate dynamic programming (ADP) could be employed. As discussed in Chapter 3, ADP combines traditional dynamic programming with Monte Carlo sampling and response surface approximation strategies to overcome the curse of dimensionality. The idea is that instead of exhaustively searching through all possible states, decisions, and information signals, ADP samples possible paths through a scenario tree to construct an approximation of the value function which can then be used to make optimal decisions for any possible state. This emerging method could be employed along with the EPPA model to create a full-fledged real-world decision-support model that would be used to explore a wide range of policy and investment decisions.

The framework developed here could also be applied to a number of other applications. One example is the study of climate impacts and investments in adaptation strategies (for example, sea walls to guard against sea level rise). This topic seems a natural fit for this framework as the impacts of a changing climate are particularly uncertain, making decisions about what adaptation strategies to pursue very difficult. This framework could also explore the balance between adaptation and mitigation strategies in the face of climate uncertainty. Ultimately, there is a wide range of problems that could be pursued using the framework developed in this dissertation.

Another future research opportunity presented by this work is a game theoretic approach to near-term investment strategies in the presence of an explicit or implicit policy "safety valve". This work shows that explicit inclusion of a safety valve in an emissions reduction policy leads to less near-term investment, increasing the probability that the safety valve will be triggered and the cap will be exceeded. One might expect that this type of behavior to also hold to some extent in the presence of an *implicit* safety valve—the prospect of the government loosening or canceling the policy in face of high costs. Investors know that the government will not enforce a policy regardless of the cost, and therefore may pursue less aggressive action with the

expectation of being "bailed out". Such an environment is amenable to game theory analysis, which could be pursued in future work.

8.6 Final Thoughts

The analytic framework and results presented in this dissertation demonstrate the importance of considering uncertainty in decision-making and highlight some of the most important issues worthy of further study. The results in this dissertation are not intended to stand alone as prescriptions, but to help decision makers think through investment decisions and policy in the context of uncertainty. It is the sincere hope of the author that this dissertation provides an analytic basis for real decision-maker discussions regarding not only electric power sector investments and policies, but also discussions and decisions involving other long-lived, complex socio-technical systems that would benefit from an economy-wide modeling approach that represents sequential decisions under uncertainty.

References

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.

Arndt, C., S. Robinson, and F. Tarp (2001). Parameter Estimation for a Computable General Equilibrium Model: A Maximum Entropy Approach. *Economic Modeling*, **19**(3): 375-398.

Arrow, K. (1962). The Economic Implications of Learning-by-Doing. *Review of Economic Studies,* **29**(3): 155–173.

Arrow, K.J. and G. Debreu (1954). Existence of an Equilibrium for a Competitive Economy. *Econometrica*, **22**(3): 265-290.

Arthur, W.B. (1989). Competing Technologies: Increasing Returns and Lock-In by Historical Events. *Economic Journal*, **99**(1): 116-131.

Azar, C. and H. Dowlatabadi. (1999). A Review of Technical Change in Assessment of Climate Policy. *Annual Review of Energy and the Environment,* **24**: 513-544.

Babiker, M.H., G.E. Metcalf, and J.M. Reilly (2003). Tax Distortions and Global Climate Policy. *Journal of Environmental Economics and Management,* **46**: 269-287.

Baker, E., H. Chon and J. Keisler (2009). Carbon Capture and Storage: Combining Economic Analysis with Expert Elicitations to Inform Climate Policy. *Climatic Change,* **96**(3): 379–408.

Balistreri, E.J., C.A. McDaniel, and E.V. Wong (2003). An Estimation of US Industry-Level Capital-Labor Substitution Elasticities: Support for Cobb–Douglas. *The North American Journal of Economics and Finance*, **14**(3): 343-356.

Basler, J.T. (2006). Optimal Portfolio Rebalancing: An Approximate Dynamic Programming Approach. Senior Thesis: Department of Operations Research and Financial Engineering, Princeton University.

Baumol, W.J. and W.E. Oates (1988). *The Theory of Environmental Policy*. Second Edition. Cambridge, UK: Cambridge University Press.

Behrens, C.E. and C. Glover (2012). U.S. Energy: Overview and Key Statistics. *Congressional Research Service.* Washington D.C.

Bellman, R.E. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University Press. Republished 2003: Dover.

Bergerson, J.A. and L.B. Lave. (2007). Baseload Coal Investment Decisions under Uncertain Carbon Legislation. *Environmental Science and Technology,* **41**(10): 3431-3436.

Berglund, C. and P. Soderholm. (2006). Modeling Technical Change in Energy System Analysis: Analyzing the Introduction of Learning-by-Doing in Bottom-up Energy Models. *Energy Policy,* **34**(12): 1344-1356.

Bertsekas, D. (2007). *Dynamic Programming and Optimal Control.* Vol. I and II. 3rd Edition. Belmont, MA: Athena Scientific.

Bertsekas, D. and J. Tsitsiklis (1996). *Neuro-Dynamic Programming.* Belmont, MA: Athena Scientific, 512pp.

Birge, J.R. and F. Louveaux (1997). *Introduction to Stochastic Programming.* New York: Springer.

Blanford, G.J. (2009). R&D Investment Strategy for Climate Change. *Energy Economics,* **31**: S27-S36, Supplement 1.

Böhringer, C., A. Löschel, U. Moslener, and T. Rutherford (2009). EU Climate Policy up to 2020: An Economic Impact Assessment. *Energy Economics*, **31**(2):S295-S305.

Böhringer, C., H. Koschel, and U. Moslener (2008). Efficiency Losses from Overlapping Regulation of EU Carbon Emissions. *Journal of Regulatory Economics*, **33**(3): 299-317.

Böhringer, C., T. Hoffman, and C. M. de Lara Peñate (2006). The Efficiency Costs of Separating Carbon Markets under the EU Emissions Trading Scheme: A Quantitative Assessment for Germany. *Energy Economics,* **28**(1): 44-61

Böhringer, C. and A. Löschel (2006). Promoting Renewable Energy in Europe: A Hybrid Computable General Equilibrium Approach. *The Energy Journal*, **27**(2): 135-150.

Botterud, A., M.D. Ilic, and I. Wangernsteen. (2005). Optimal Investments in Power Generation under Centralized and Decentralized Decision Making. *IEEE Transactions on Power Systems,* **20**(1): 254-263.

Bosetti, V., M. Catenacci, G. Fiorese, and E. Verdoloni (2012). The Future Prospect of PV and CSP Solar Technologies: An Expert Elicitation Survey. *Energy Policy*, **49**:308-317.

Bosetti, V., C. Carraro, and E. Massetti (2009). Banking Permits: Economic Efficiency and Distributional Effects. *Journal of Policy Modeling*, **31**(3):381-403.

Bossetti, V. and M. Tavoni. (2009). Uncertain R&D, Backstop Technology, and GHGs Stabilization. *Energy Economics,* **31**: S18-S26.

Bosetti, V., C. Carraro, M. Galeotti, E. Massetti, and M. Tavoni. (2006). WITCH: A World Induced Technical Change Hybrid Model. *The Energy Journal*, Special Issue on Hybrid Modeling of Energy-Environment Policies: Reconciling Bottom-up and Top-down**.**

Burtraw, D. and K. Palmer. (2006). Dynamic Adjustment to Incentive-Based Policy to Improve Efficiency and Performance. Washington, DC: Resources for the Future.

Chan, G., L.D. Anadon, M. Chan, and A. Lee (2010). Expert Elicitation of Cost, Performance, and RD&D Budgets for Coal Power with CCS. *Energy Procedia* [Internet]. **4**:2685-2692.

Clarke, L., J. Edmonds, V. Krey, R. Richels, S. Rose, and M. Tavoni (2009). International climate policy architectures: Overview of the EMF 22 International Scenarios. *Energy Economics*, **31**:S64-S81.

Clarke, L., J. Weyant, and J. Edmonds. (2008). On the Sources of Technological Change: What Do the Models Assume? *Energy Economics,* **30**(2): 409-424.

Clarke, L., J. Edmonds, H. Jacoby, H. Pitcher, J. Reilly, R. Richels (2007). Scenarios of Greenhouse Gas Emissions and Atmospheric Concentrations, Sub-report 2.1A of Synthesis and Assessment Product 2.1 by the U.S. Climate Change Science Program and the Subcommittee on Global Change Research, Department of Energy, Office of Biological and Environmental Research, Washington, DC

Clemen, R.T. and R.L. Winkler (1999). Combining Probability Distributions from Experts in Risk Analysis. *Risk Analysis*, **19**(2): 187-203.

Colpier, U.C. and D. Cornland (2002). The Economics of the Combined Cycle Gas Turbine-An Experience Curve Analysis. *Energy Policy,* **30**(4): 309–316.

Crost, B. and C. Traeger. (2010). Risk and Aversion in the Integrated Assessment of Climate Change. CUDARE Working Paper No. 1104.

Dantzig, G.B. and P. Wolfe. (1960). Decomposition Principle for Linear Programs. *Operations Research*, **8**(1): 101–111.

deWeck, O.L., D. Roos, and C.L. Magee (2011). *Engineering Systems: Meeting Human Needs in a Complex Technological World*. Cambridge, MA: MIT Press.

Dimaranan, B. and R.McDougall (2002). Global Trade, Assistance, and Production: The GTAP 5 Data Base. Center for Global Trade Analysis, Purdue University, West Lafayette, Indiana.

Dobesova, K., J. Apt and L.B. Lave (2005). Are Renewables Portfolio Standards Cost-Effective Emission Abatement Policy? Carnegie Mellon Electricity Industry Center Working Paper CEIC-04-06, Pittsburgh, PA.

Dosi, G. (1988). The Nature of the Innovative Process. In: Dosi, G., C. Freeman, C. *et al.* (Eds.), *Technical Change and Economic Theory*, Printer Ed., London

Edmonds, J., M. Wise, and D. W. Barns (1995). Carbon Coalitions: The Cost and Effectiveness of Energy Agreements to Alter Trajectories of Atmospheric Carbon Dioxide Emissions. *Energy Policy,* **23**(4/5): 309-335.

Edmonds, J. and J.M. Reilly (1985). *Global Energy: Assessing the Future.* New York: Oxford University Press.

EIA [Energy Information Administration] (2013a). Annual Energy Outlook 2013: With Projections to 2040. U.S. Department of Energy, Washington, DC.

EIA (2013b). Electric Power Annual 2011. U.S. Department of Energy, Washington, DC.

EIA (2013c). Electric Power Monthly with Data for June 2013. U.S. Department of Energy, Washington, DC.

EIA (2012). Annual Energy Review 2011. U.S. Department of Energy, Washington, DC.

EIA (2009). Energy Market and Economic Impacts of H.R. 2454, the American Clean Energy and Security Act of 2009. U.S. Department of Energy, Washington, DC.

Ellerman, A. D., M. D. Webster, J. Parsons, H. D. Jacoby, M. McGuinness (2008). *Cap-and-Trade: Contributions to the Design of a U.S. Greenhouse Gas Program*. MIT Center for Energy and Environmental Policy Research, Cambridge, MA

Ereira, E.C. (2010). Assessing Early Investments in Low Carbon Technologies under Uncertainty: The Case of Carbon Capture and Storage. Master of Science Thesis, Technology and Policy Program, MIT.

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. and R.G. Newell (2004). Environmental and Technology Policies for Climate Change and Renewable Energy. RFF Discussion Paper 04-05. Washington, DC: Resources for the Future.

Fisher, B.S., N. Nakicenovic, K. Alfsen, J. Corfee Morlot, F. de la Chesnaye, J.-Ch. Hourcade, K. Jiang, M. Kainuma, E. La Rovere, A. Matysek, A. Rana, K. Riahi, R. Richels, S. Rose, D. van Vuuren, R. Warren (2007). Issues Related to Mitigation in the Long Term Context. In: *Climate Change 2007: Mitigation*. Contribution of Working Group III to the Fourth Assessment Report of the Inter-governmental Panel on Climate Change, B. Metz, O.R. Davidson, P.R. Bosch, R. Dave, L.A. Meyer (eds.), Cambridge, U.K.: Cambridge University Press.

Genest, C. and J.V. Zidek (1986). Combining Probability Distributions: a Critique and Annotated Bibliography. *Statistical Science* **1**(1): 114-148.

George, A. and W.B. Powell (2006). Adaptive Stepsizes for Recursive Estimation with Applications in Approximate Dynamic Programming. *Machine Learning,* **65**(1): 167-198.

Gerst, M., R. Howarth, and M. Borsuk. (2010). Accounting for the risk of extreme outcomes in an integrated assessment of climate change. *Energy Policy,* **38**(8): 4540-4548.

Godfrey, G. and W.B. Powell (2002). An Adaptive, Dynamic Programming Algorithm for Stochastic Resource Allocation Problems I: Single Period Travel Times. *Transportation Science,* **36**(1): 21-39.

Goulder, L. H. (1995). Environmental Taxation and the 'Double Dividend': A Reader's Guide. *International Tax and Public Finance*, **2**(2): 157-183.

Grubb, M., J. Köhler, and D. Anderson. (2002). Induced Technical Change in Energy and Environmental Modeling: Analytic Approaches and Policy Implications. *Annual Review of Energy and the Environment,* **27**: 271-308.

Grubler, A., and A. Gritsevskii. (1997). A Model of Endogenous Technological Change Through Uncertain Returns on Learning (R&D and investments). Conf. Induc. Technol. Change Environ. International Institute of Applied System Analysis, Laxenburg, Austria.

Gurgel, A., S. Paltsev, J. Reilly and G. Metcalf (2007). U.S. Greenhouse Gas Cap-and-Trade Proposals: Application of a Forward-Looking Computable General Equilibrium Model. MIT Joint Program on the Science and Policy of Global Change *Report 150*. (Available at: http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt150.pdf).

Haikel, K.M. (2009). A Game Theoretic Model for Generation Capacity Adequacy in Electricity Markets: A Comparison Between Investment Incentive Mechanisms. GATE Groupe d'Analyse et de Théorie Économique Working Papers.

Hammitt, J. K., R. A. Lempert, and M. E. Schlesinger (1992). A Sequential-Decision Strategy for Abating Climate Change. *Nature,* **357**(6376): 315-318.

Hertel, T. (1997). *Global Trade Analysis: Modeling and Applications*. Cambridge, UK: Cambridge University Press.

Ibenholt, K. (2002). Explaining Learning Curves for Wind Power. *Energy Policy,* **30**(13): 1181-1189.

IEA [International Energy Agency] (2013). Key World Energy Statistics. International Energy Ageny, Paris, France. Available at: http://www.iea.org/publications/freepublications/publication/kwes.pdf

IPCC [Intergovernmental Panel on Climate Change] (2007). Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK: Cambridge University Press.

Jacoby, H.D., J.M. Reilly, J.R. McFarland and S. Paltsev (2006). Technology and Technical Change in the MIT EPPA Model. *Energy Economics*, **28**(5-6): 610-631.

Jacoby, H.D., and A.D. Ellerman (2004). The Safety Valve and Climate Policy. *Energy Policy,* **32**(4): 481-491.

Jacoby, H.D., R.S. Eckhaus, A.D. Ellerman, R.G. Prinn, D.M. Reiner and Z. Yang (1997). CO_2 Emissions Limits: Economic Adjustments and the Distribution of Burdens. *The Energy Journal,* **18**(3): 31-58.

Janetos, A.C., L. Clarke, B. Collins, K. Ebi, J. Edmonds, I. Foster, H.D. Jacoby, K. Judd, R. Leung, R.Newell, D. Ojima, G. Pugh, A. Sanstad, P. Schultz, R. Stevens, J. Weyant, T. Wilbanks (2009). Science Challenges and Future Directions: Climate Change Integrated Assessment Research. Report from the U.S. Department of Energy, Office of Science, Office of Biological and Environmental Research, Workshop on Integrated Assessment, November 2008. PNNL-18417.

Kall, P. and S.W. Wallace (1994). *Stochastic Programming.* Hoboken, NJ: John Wiley & Sons, Inc.

Karkatsouli, I. (2013). The Role of Wind Generation in European Power Sector Decarbonization: A General Equilibrium Analysis. Master of Science Thesis, Technology and Policy Program, MIT.

Kelly, D. and C. Kolstad. (1999). Bayesian Learning, Growth, and Pollution. *Journal of Economic Dynamics and Control,* **23**(4): 491-518.

Klein, A., J. Bouchard, and S. Goutier (2008). Generation Capacity Expansion under Long-Term Uncertainties in the U.S. Electric Market. United States Association of Energy Economics Conference Paper.

Klepper G. and S. Peterson (2004). The EU Emissions Trading Scheme: Allowances Prices, Trade Flows, Competitive Effects. *Environmental Policy and Governance,* **14**(4): 201-218.

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.

Kypreos, S. and O. Bahn. (2003). A MERGE Model with Endogenous Technological Progress. *Environmental Modeling and Assessment,* **8**(3): 249-259.

Kypreos, S. and L. Barreto (2000). ERIS: A Model Prototype with Endogenous Technological Change. *International Journal of Global Energy Issues,* **14**(1-4): 374-396.

Lechon, Y., H. Cabal, M. Varela, R. Saez, C. Eherer, M. Baumann, J. Duweke, T. Hamacher, and G.C. Tosato. (2005). A Global Energy Model with Fusion. *Fusion Engineering and Design,* **75-79**: 1141-1144.

Leiserowitz, A., N. Smith, and J.R. Marlon (2010). Americans' Knowledge of Climate Change. Yale Project on Climate Change Communication, Yale University: New Haven, CT.

Lemaréchal, C. (2001). Lagrangian Relaxation. In: M. Jünger and D. Naddef (eds.), *Computational Combinatorial Optimization: Optimal or Provably Near-Optimal Solutions.* New York: Springer.

Lemoine, D. and C. Traeger. (2011). Tipping Points and Ambiguity in the Integrated Assessment of Climate Change. NBER Environmental and Energy Economics Summer Institute Paper.

Little, J. D. C. (1970). Models and Managers: The Concept of a Decision Calculus. *Management Science*, **16**(8): B466-B485.

Löschel, A. and V.M. Otto (2009). Technological Uncertainty and Cost Effectiveness of CO2 Emission Reduction. *Energy Economics*, **31**: S4-S17.

Loulou, R., G. Goldstein, and K. Noble. (2004). Documentation for the MARKAL Family of Models. Energy Technology Systems Analysis Programme. Available at: http://www.etsap.org/tools.htm

Manne, A.S. and R.G. Richels (1995a). The Greenhouse Debate: Economic Efficiency, Burden Sharing and Hedging Strategies. *The Energy Journal,* **16**(4): 1-37.

Manne, A.S., R. Mendelsohn, and R.G. Richels (1995b). MERGE: A Model for Evaluating Regional and Global Effects of GHG Reduction Policies. *Energy Policy,* **23**(1): 17-34.

Manne, A.S. and R.G. Richels (1994). The Costs of Stabilizing Global CO_2 Emissions: A Probabilistic Analysis Based on Expert Judgment. *The Energy Journal,* **15**(1): 31-56.

Mathiesen, L. (1985). Computation of Economic Equilibrium by a Sequence of Linear Complementarity Problems. *Mathematical Programming Study*, **23**: 144-162.

Mattsson, N. (2002). Introducing Uncertain Learning in an Energy System Model: A Pilot Study Using GENIE. *International Journal of Global Energy Issues,* **18**(2/3/4): 253-265.

Mattsson, N. and C.O. Wene. (1997). Assessing New Energy Technologies Using an Energy System Model with Endogenized Experience Curves. *International Journal of Energy Research,* **21**(4): 385-393.

McDonald, A. and L. Schrattenholzer (2001). Learning Rates for Energy Technologies. *Energy Policy,* **29**(4): 255-261.

McFarland, J.R. and H.J. Herzog (2006). Incorporating Carbon Capture and Storage Technologies in Integrated Assessment Models. *Energy Economics*, **28**(5-6):632-652.

McFarland, J.R., J.M. Reilly and H.J. Herzog, (2004). Representing Energy Technologies in Top-Down Economic Models Using Bottom-Up Information. *Energy Economics*, **26**(4): 685-707.

Messner, S. (1997). Endogenized Technological Learning in an Energy Systems Model. *Journal of Evolutionary Economics,* **7**(3): 291-313.

Messner, S., A. Golodnikov, and A. Gritsevskii. (1995). A Stochastic Version of the Dynamic Linear Programming Model MESSAGE III. *Energy,* **21**(9): 775-784.

Metcalf, G.E., S. Paltsev, J.M. Reilly, H.D. Jacoby and J. Holak (2008). Analysis of U.S. Greenhouse Gas Tax Proposals. MIT Joint Program on the Science and Policy of Global Change *Report 160*, (Available at: http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt160.pdf).

Metcalf, G.E., M. Babiker and J.M. Reilly (2004). A Note on Weak Double Dividends. *Topics in Economic Analysis and Policy*, **4**(1): Article 2.

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.

Morris, J.F., S. Paltsev, and J.M. Reilly (2012). Marginal Abatement Costs and Marginal Welfare Costs for Greenhouse Gas Emissions Reductions: Results from the EPPA Model. *Environmental Modeling and Assessment*, **17**(4): 325-336.

Morris, J.F., J.M. Reilly and S. Paltsev (2010). Combining a Renewable Portfolio Standard with a Cap-and-Trade Policy: A General Equilibrium Analysis. MIT Joint Program on the Science and Policy of Global Change *Report 187*, (Available at: http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt187.pdf).

Morris, S.C., G.A. Goldstein, and V.M. Fthenakis (2002). NEMS and MARKAL-MACRO Models for Energy-Environmental-Economic Analysis: A comparison of the Electricity and Carbon Reduction Projections. *Environmental Modeling and Assessment* **7**(3): 207-216.

Moss, R., M. Babiker, S. Brinkman, E. Calvo, T. Carter, J. Edmonds, I. Elgizouli, S. Emori, L. Erda, K. Hibbard, R. Jones, M. Kainuma, J.Kelleher, J. Francois Lamarque, M. Manning, B. Matthews, J. Meehl, L. Meyer, J. Mitchell, N. Nakicenovic, B. O'Neill, R. Pichs, K. Riahi, S. Rose, P. Runci, R. Stouffer, D. van Vuuren, J. Weyant, T. Wilbanks, J. Pascal van Ypersele, M. Zurek, F. Birol, P. Bosch, O. Boucher, J. Feddema, A. Garg, A. Gaye,

M. Ibarraran, E. La Rovere, B. Metz, S. Nishioka, H. Pitcher, D. Shindell, P.R. Shukla, A. Snidvongs, P. Thorton, V. Vilariño (2007). Towards New Scenarios for Analysis of Emissions, Climate Change, Impacts, and Response Strategies, IPCC Expert Meeting Report, Noordwijkerhout, The Netherlands. Available at: http://ipcc-data.org/docs/ar5scenarios/IPCC_Final_Draft_Meeting_Report_3May08.pdf.

Mowery, D.C. and A.A. Ziedonis (2001). The Geographic Reach of Market and Non-Market Channels of Technology Transfer: Comparing Citations And Licenses Of University Patents. Working Paper 8568 National Bureau Of Economic Research, Cambridge.

Murphy, F. H. and Y. Smeers. (2005). Generation Capacity Expansion in Imperfectly Competitive Restructured Electricity Markets. *Operations Research,* **53**(4): 646-661.

National Research Council (2007). Prospective Evaluation of Applied Energy Research and Development at DOE (Phase Two). The National Academies Press, Washington D.C.

Neij, L. (1997). Use of Experience Curves to Analyse the Prospects for Diffusion and Adoption of Renewable Energy Technology. *Energy Policy* **23**(13): 1099–1107.

Newell, R.G. and W.A. Pizer (2003). Regulating Stock Externalities under Uncertainty. *Journal of Environmental Economics and Management,* **45:** 416-432.

Nogee, A., J. Deyette and S. Clemmer (2007). The Projected Impacts of a National Renewable Portfolio Standard. *The Electricity Journal*, **20**(4): 33-47.

Nordhaus, W.D. and D. Popp. (1997). What is the Value of Scientific Knowledge? An Application to Global Warming Using the PRICE Model. *The Energy Journal,* **18**(1): 1-45.

Nordhaus, W. D. and Z. Yang (1996). A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies. *American Economic Review*, **86**(4): 741-765.

Nordhaus, W.D. (1992). The 'DICE' Model: Background and Structure of a Dynamic Integrated Climate-Economy Model of the Economics of Global Warming. Cowles Foundation Discussion Papers 1009, Cowles Foundation for Research in Economics, Yale University.

Ormoneit, D. and S. Sen (2002). Kernel-Based Reinforcement Learning. *Machine Learning,* **49**: 161-178.

Osouf, N. (2007). The Potential for a Nuclear Renaissance: The Development of Nuclear Power under Climate Change Mitigation Policies. Master of Science Thesis, Technology and Policy Program and Department of Nuclear Science and Engineering, MIT.

Otto, V.M., and J.M. Reilly (2008). Directed Technical Change and the Adoption of CO2 Abatement Technology: The Case of CO2 Capture and Storage. *Energy Economics*, **30**(6): 2879-2898.

Palmer, K. and D. Burtraw (2005). Cost-Effectiveness of Renewable Electricity Policies. RFF Discussion Paper 05-01. Washington, D.C., Resources for the Future.

Paltsev, S., H.D. Jacoby, J.M. Reilly, Q.J. Ejaz, J.F. Morris, F. O'Sullivan, S. Rausch, N. Winchester and O. Kragha (2011). The Future of U.S. Natural Gas Production, Use, and Trade. *Energy Policy*, **39**(9): 5309-5321.

Paltsev, S., J. M. Reilly, H.D. Jacoby, and J.F Morris (2009). The Cost of Climate Policy in the United States. MIT Joint Program on the Science and Policy of Global Change *Report 173*, (Available at: http://globalchange.mit.edu/pubs/abstract.php?publication_id=972).

Paltsev, S., J.M. Reilly, H.D. Jacoby, A.C. Gurgel, G.E. Metcalf, A.P. Sokolov, and J.F. Holak (2008). Assessment of US GHG Cap-and-Trade Proposals. *Climate Policy*, **8**(4):395-420.

Paltsev, S., J.M. Reilly, H.D. Jacoby, R.S. Eckaus, J. McFarland, M. Sarofim, M. Asadoorian, M. Babiker (2005). The MIT Emissions Prediction and Policy Analysis (EPPA) Model: Version 4. MIT Joint Program on the Science and Policy of Global Change *Report 125*, (Available at: http://web.mit.edu/globalchange/www/reports.html#pubs).

Parpas, P. and M. Webster (2012). A stochastic minimum principle and a mesh-free method for stochastic optimal control. *Revise-and-resubmit*.

PCAST (1997). Report to the President on Federal Energy Research and Development for the Challenges of the Twenty-First Century. Washington DC, PCAST.

Pérez-Arriaga, I.J. (2013). Challenges in Power Sector Regulation. In: I. J. Pérez-Arriaga (ed.), *Regulation of the Power Sector.* New York: Springer.

Pérez-Arriaga, I. J., H. Rudnick, and M. Rivier (2008). Electric Energy Systems-An Overview. In: A. Gómez-Expósito, A.J. Conejo, and C. Canizares (eds.), *Electric Energy Systems: Analysis and Operation.* Boca Raton: CRC Press.

Phelps, E. S. (1963). Substitution, Fixed Proportions, Growth and Distribution. *International Economic Review*, **4**(3): 265-288.

Pizer, W.A. (2005). Climate Policy Design under Uncertainty (Discussion Paper).Washington D.C.: Resources for the Future.

Pizer, W.A. (2002) Combining Price and Quantity Controls to Mitigate Global Climate Change. *Journal of Public Economics*, **85**(3): 409-434.

Pizer, W.A. (1997): Prices vs. Quantities Revisited: The Case of Climate Change. *RFF Discussion Paper 98-02.* Washington, D.C., Resources for the Future

Popp, D. (2006). ENTICE-BR: Backstop Technology in the ENTICE Model of Climate Change. *Energy Economics,* **28**(2): 188-222.

Popp, D. (2004). ENTICE: Endogenous Technological Change in the DICE Model of Global Warming. *Journal of Environmental Economics and Management,* **48**(1): 742-768.

Powell, W.B., A. George, H. Simão, W. Scott, A. Lamont and J. Stewart. (2012). SMART: A Stochastic Multiscale Model for the Analysis of Energy Resources, Technology and Policy. *INFORMS Journal on Computing,* **24**(4): 665-682

Powell, W.B. (2011). *Approximate Dynamic Programming: Solving the Curses of Dimensionality.* 2nd Ed. Hoboken, NJ: John Wiley & Sons, Inc.

Puterman, M.L. (2005). *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* Hoboken, NJ: John Wiley & Sons, Inc.

Pyatt, G., and J.I. Round (1985). Social Accounting Matrices: A Basis for Planning. The World Bank: Washington D.C.

Rausch, S. and M. Mowers (2012). Distributional and Efficiency Impacts of Clean and Renewable Energy Standards for Electricity. MIT Joint Program on the Science and Policy of Global Change *Report 225*, (Available at: http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt225.pdf).

Rausch, S., G.E. Metcalf, J.M. Reilly and S. Paltsev (2010). Distributional Implications of Alternative U.S. Greenhouse Gas Control Measures. MIT Joint Program on the Science and Policy of Global Change *Report 185*, (Available at: http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt185.pdf).

Reilly J.M., J. Melillo, C. Yongxia, D. Kicklighter, A. Gurgel, S. Paltsev, T. Cronin, A.P. Sokolov, and A. Schlosser (2012). Using Land To Mitigate Climate Change: Hitting the Target, Recognizing the Trade-offs. *Environmental Science and Technology*, **46**(11): 5672–5679 (Supporting Information).

Reilly, J.M. and S. Paltsev (2006). European Greenhouse Gas Emissions Trading: A System in Transition. In: C. de Miguel, X. Labandeira, and B. Manzano (eds.), *Economic Modeling of Climate Change and Energy Policies.* Northampton, MA: Edward Elgar Publishing.

Reilly, J.M., R.G. Prinn, J. Harnisch, J. Fitzmaurice, H.D. Jacoby, D. Kicklighter, J. Melillo, P.H. Stone, A.P. Sokolov, and C. Wang (1999). Multi-Gas Assessment of the Kyoto Protocol. *Nature,* **401**: 549-555.

Reilly, J.M., J. Edmonds, R. Gardner and A. Brenkert (1987). Monte Carlo Analysis of the IEA/ORAU Energy/Carbon Emissions Model. *The Energy Journal,* **8**(3): 1-29.

Richels, R.G. and G.J. Blanford. (2008). The Value of Technological Advance in Decarbonizing the U.S. Economy. *Energy Economics,* **30**(6): 2930-2946.

Roberts, M. J. and M. Spence (1976). Effluent Charges and Licenses under Uncertainty. *Journal of Public Economics*, **5**(3-4): 193-208.

Rockafellar, R.T. and R. J.-B. Wets (1991). Scenarios and Policy Aggregation in Optimization under Uncertainty. *Mathematics of Operations Research*, **16**(1):119–147.

Ross, M. (2008). Documentation of the Applied Dynamic Analysis of the Global Economy (ADAGE) Model. Research Triangle Institute Working Paper 08_01, September 2008. Available at: http://www.rti.org/pubs/adage-model-doc_ross_apr07.pdf

Rutherford, T.F. (1999). Applied General Equilibrium Modeling with MPSGE as a GAMS Subsystem: An Overview of the Modeling Framework and Syntax. *Computational Economics*, **14**: 1-46.

Rutherford, T.F. (1995). Extension of GAMS for Complementarity Problems Arising in Applied Economic Analysis. *Journal of Economic Dynamics and Control*, **19**(8): 1299-1324.

Sanchez, J.J., E. Centeno, and J. Barquin (2012). System Dynamics Modeling for Electricity Generation Expansion Analysis. Institute for Research and Technology, Universidad Pontificia Comillas, Working Paper.

Sanchez, J.J. (2008). Strategic Analysis of the Long-Term Planning of Electric Generation Capacity in Liberalized Electricity Markets. European Ph.D. Thesis submitted to Universidad Pontificia Comillas de Madrid.

Santen, N.R. (2012). Technology Investment Decisions under Uncertainty: A New Modeling Framework for the Electric Power Sector. Massachusetts Institute of Technology Doctoral Dissertation, Engineering Systems Division.

Scott, M.J., R.D. Sands, J. Edmonds, A.M. Liebetrau, and D.W. Engel. (1999). Uncertainty in Integrated Assessment Models: Modeling with MiniCAM 1.0. *Energy Policy,* **27**(14): 597.

Seebregts. A.J., T. Kram, G.J. Schaeffer, A. Stoffer, S. Kypreos, L. Barreto, S. Messner, and L. Schrattenholzer (1999). Endogenous Technological Change in Energy System Models, Synthesis of Experience with ERIS, MARKAL, and MESSAGE. Report ECN-C—99-025 of the Paul Scherrer Institute and the International Institute for Applied Systems Analysis.

Short, W., N. Blair, P. Sullivan, and T. Mai. (2009). ReEDS Model Documentation: Base Case Data and Model Description. National Renewable Energy Laboratory

Shoven, J. and J.Whalley (1984). Applied General-Equilibrium Models of Taxation and International Trade: An Introduction and Survey. *Journal of Economic Literature*, **22**(3): 1007-1051.

Smith, A.E. (2007). CRA International's MRN-NEEM Intergrated Model for Analysis of US Greenhouse Gas Policies. Charles River Associates, Boston, MA.

Sokolov, A.P., P.H. Stone, C.E. Forest, R. Prinn, M.C. Sarofim, M. Webster, S. Paltsev, C.A. Schlosser, D. Kicklighter, S. Dutkiewicz, J.M. Reilly, C. Wang, B. Felzer, J.M. Melillo, and H.D. Jacoby (2009). Probabilistic Forecast for 21st Century Climate Based on Uncertainties in Emissions (Without Policy) and Climate Parameters. *Journal of Climate*, **22**(19): 5175-5204.

Sokolov, A.P., C.A. Schlosser, S. Dutkiewicz, S. Paltsev, D.W. Kicklighter, H.D. Jacoby, R.G. Prinn, C.E. Forest, J.M. Reilly, C. Wang, B. Felzer, M.C. Sarofim, J. Scott, P.H. Stone, J.M. Melillo, and J. Cohen (2005). The MIT Integrated Global System Model (IGSM) Version 2: Model Description and Baseline Evaluation. MIT Joint Program on the Science and Policy of Global Change *Report 124*, 40 pp. (Available on line at http://web.mit.edu/globalchange/www/MITJPSPGC_Rpt124.pdf)

Staël von Holstein, C.-A. S. and J.E. Matheson (1979). *A Manual for Encoding Probability Distributions*. Palo Alto, CA: SRI International.

Stavins, R.N. (1997). Policy Instruments for Climate Change: How Can National Governments Address a Global Problem? *The University of Chicago Legal Forum,* volume 1997: 293-329.

Sue Wing, I. (2004). Computable General Equilibrium Models and Their Use in Economy-Wide Policy Analysis. Joint Program on the Science and Policy of Global Change *Technical Note 6*, (Available at: http://globalchange.mit.edu/files/document/MITJPSPGC_TechNote6.pdf).

Tietenberg, T.H. (1990). Economic Instruments for Environmental Regulation. *Oxford Review of Economic Policy*, **6**(1): 17-33.

Tversky, A. and D. Kahneman, (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, **185**(4157): 1124-1131.

US EPA [Environmental Protection Agency] (2010). Documentation for EPA Base Case v.4.10 Using the Integrated Planning Model. EPA Report #430R10010, August 2010. Available at: http://www.epa.gov/airmarkt/progsregs/epa-ipm/docs/v410/Chapter1.pdf.

US EPA (2009a). Endangerment and Cause or Contribute Findings for Greenhouse Gases under Section 202(a) of the Clean Air Act. *Federal Register* Docket ID No. EPA-HQ-OAR-2009-0171, December 15, 2009, Washington DC.

US EPA (2009b). EPA Analysis of the American Clean Energy and Security Act of 2009 H.R. 2454 in the 111th Congress. Environmental Protection Agency, Washington D.C. Available at: http://www.epa.gov/climatechange/Downloads/EPAactivities/HR2454_Analysis.pdf

USGCRP [U.S. Global Change Research Program] (2009). Global Climate Change Impacts in the United States. New York: Cambridge University Press.

U.S. Government (2013). Technical Update of Social Cost of Carbon For Regulatory Impact Analysis. Technical Support Document under Executive Order 12866. Interagency Working Group on Social Cost of Carbon. Washington, D.C.

van der Zwaan, B. and A. Rabl (2003). Prospects for PV: A Learning Curve Analysis. *Solar,* **74**(1): 19–31.

van der Zwaan, B., R. Gerlagh, G. Klassen, and L. Schrattenholzer. (2002). Endogenous Technological Change in Climate Change Modeling. *Energy Economics,* **24**(1): 1-19.

Virdis, M., C. Wene, and H. Nilsson (2003). Creating Markets for Energy Technologies. International Energy Agency, Paris.

Wang, C., and R. Prinn (2010). Potential Climatic Impacts and Reliability of Very Large-Scale Wind Farms. *Atmospheric Chemistry and Physics*, **10**(4): 2053-2061.

Webster, M.D., N.R. Santen, and P. Parpas. (2012a). An Approximate Dynamic Programming Framework for Modeling Global Climate Policy under Decision-Dependent Uncertainty. *Computational Management Science,* **9**(3): 339-362.

Webster, M.D., A.P. Sokolov, J.M. Reilly, C. Forest, S. Paltsev, C.A. Schlosser, C. Wang, D.W. Kicklighter, M. Sarofim, J.M. Melillo, R.G. Prinn and H.D. Jacoby (2012b). Analysis of Climate Policy Targets under Uncertainty. *Climate Change,* **112**(3-4): 569-583.

Webster, M.D., I. Sue Wing, and L. Jakobovits (2010). Second-Best Instruments for Near-Term Climate Policy: Intensity Targets vs. the Safety Valve. *Journal of Environmental Economics and Management,* **59**(3): 250-259.

Webster, M.D., S. Paltsev, J. Parsons, J. Reilly, and H.D. Jacoby (2008a). Uncertainty in Greenhouse Emissions and Costs of Atmospheric Stabilization. MIT Joint Program on the Science and Policy of Global Change *Report 165*. (Available at: http://web.mit.edu/globalchange/www/MITJPSPGC_Rpt165.pdf).

Webster, M.D., L. Jakobovits, and J.Norton. (2008b). Learning about Climate Change and Implications for Near-Term Policy. *Climatic Change,* **89**(1-2): 67-85.

Webster, M.D. (2008c). Incorporating Path-Dependency into Decision Analytic Methods: An Application to Global Climate Change Policy. *Decision Analysis,* **5**(2): 60-75.

Webster, M.D. (2002). The Curious Role of Learning: Should We Wait for More Data? *The Energy Journal,* **23**(2): 97-119.

Weitzman, M.L. (1974). Prices vs. Quantities. *Review of Economic Studies*, **41**(4): 477-491.

Wigley, T.M.L., R. Richels, and J.A. Edmonds (1996). Economic and Environmental Choices in the Stabilization of Atmospheric CO_2 Concentrations. *Nature*, **379**(6562): 240-243.

Wiser, R., C. Namovicz, M. Gielecki and R. Smith (2007). The Experience with Renewable Portfolio Standards in the United States. *The Electricity Journal*, **20**(4): 8-20.

Ybema, J.R., Bos, A.J.M., and A.J. Seebregts. (1998). Incorporating the Long Term Risk for Deep Emission Reduction in Near Term CO2 Mitigation Strategies. *Journal of Hazardous Materials,* **61**(1-3): 217-227.

Yeh, S. and E. Rubin. (2007). A Centurial History of Technological Change and Learning Curves on Pulverized Coal-Fired Utility Boilers. UC Davis: Institute of Transportation Studies. Available at: http://pubs.its.ucdavis.edu/publication_detail.php?id=1108.

Yohe, G., N.G. Andronova, and M.E. Schlesinger. (2004). To Hedge or Not Against an Uncertain Climate Future? *Science,* **306**(5695): 416-417.

Zhang, X. and G. Verikios (2006). Armington Parameter Estimation for a Computable General Equilibrium Model: A Database Consistent Approach. Economics Discussion Working Papers No.06-10. Crawley, Australia: The University of Western Australia.

(This page is intentionally left blank)

Appendix A: Social Accounting Matrix

Table A.1 Social Accounting Matrix (SAM) for CGE Model: 2010 Base Year Data ($, billion)

		Sector						
		Electricity	Coal	Natural Gas	Oil	Crude Oil	Other	Household
Input	**Electricity**						155.0	100.0
	Coal	19.5					3.0	1.5
	Natural Gas	18.0					7.0	20.0
	Oil						50.0	100.0
	Crude Oil				90.0			
	Other	97.5	7.2	19.5	52.5	51.0	9245.0	7037.3
	Capital	90.0	4.8	15.0	4.5	21.0	2850.0	
	Labor	30.0	4.8	6.0	3.0	3.0	4200.0	
	Coal Resource		7.2					
	Natural Gas Resource			4.5				
	Crude Oil Resource					15.0		
	Total Production	255.0	24.0	45.0	150.0	90.0	16510.0	

(This page is intentionally left blank)

Appendix B: Survey on Policy Expectations

B.1 Actual Survey

The Survey was presented in an online format using SuveyMonkey.com. The content of the survey as it appeared is included below.


Survey on the Likelihood of Emissions Caps in the Electric Power Sector

This survey is designed to attain the perceived likelihood of different levels of CO_2 emission caps in the electric power sector. The caps are formulated as **cumulative reductions** in CO_2 emissions over the period 2015-2030.

Please assign probabilities to each of the following cap policies. Please ensure that your probabilities sum to 100%. See the graph below for guidance on the different cap levels.

Cumulative Cap (2015-2030)	Probability
0% (no cap) to <10%	0%
10% to <20%	0%
20% to <30%	0%
30% to <40%	0%
40% to <50%	0%
50% or greater	0%

The graph below shows reference electricity emissions (from EIA's Annual Energy Outlook 2012) and linear reduction paths that result in given cumulative emissions reductions from 2015-2030. The cumulative reduction is the area of the triangle between the reference path and a policy path. Of course other, non-linear emissions paths could also meet the same cumulative cap (for example, fewer reductions early on and more in later years, or vice versa). The graph also shows where the linear reduction paths end up in 2030 relative to present emissions. Regarding the ranges in the survey table above, a cumulative cap from 0% to <10% is represented by the top "wedge" (between reference and 10% cumulative cap). A cumulative cap from 10% to <20% is represented by the second "wedge" (between 10% and 20% cumulative caps), and so forth.

Electricity Sector Emissions Paths
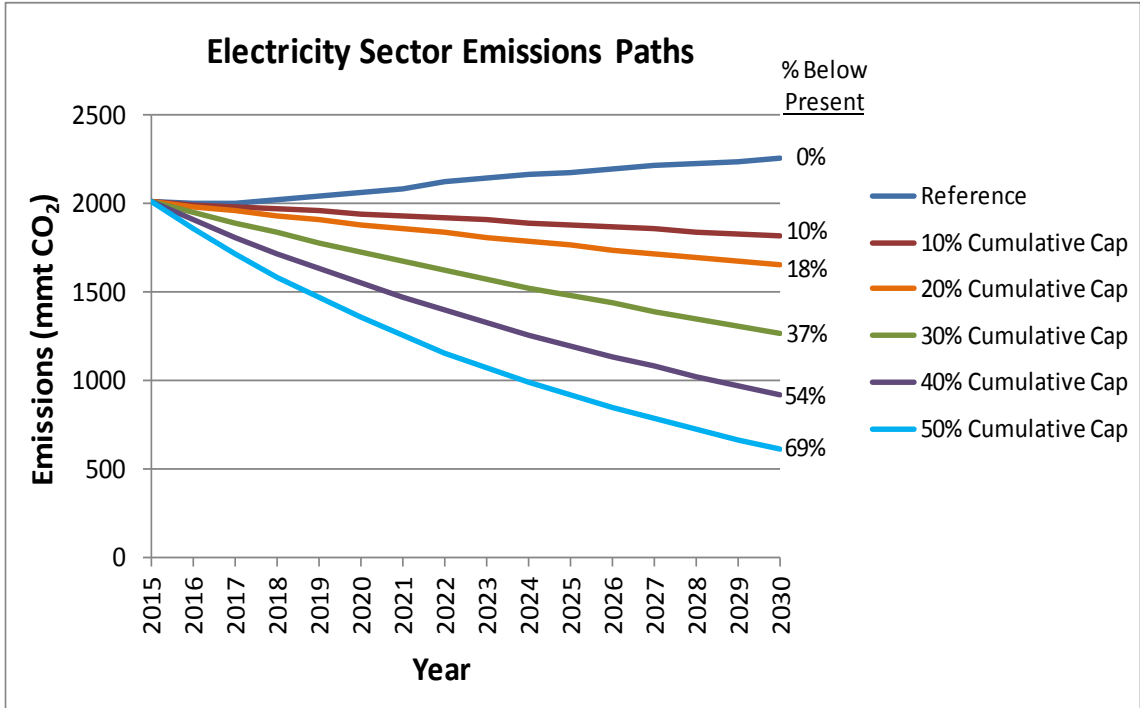
B.2 Full Survey Results

Table B.2.1 Full Survey Results

Cumulative Cap (2015-2030)	Respondents' Expected Probability of Each Policy						
	1	**2**	**3**	**4**	**5**	**6**	**7**
0% (no cap) to <10%	60%	50%	60%	35%	20%	10%	5%
10% to <20%	35%	50%	30%	25%	0%	15%	5%
20% to <30%	5%	0%	10%	15%	40%	20%	10%
30% to <40%	0%	0%	0%	10%	20%	30%	40%
40% to <50%	0%	0%	0%	10%	20%	20%	30%
50% or greater	0%	0%	0%	5%	0%	5%	10%

The survey inquired about six possible policy ranges for emissions reductions. However, this research focused on three possible policies: no policy, -20% cap, and -40% cap. The survey results were therefore distributed into those three categories. This was done using the following equations:

Expected Probability of No Policy = (probability assigned to "0% (no cap) to <10%" policy) + ½(probability assigned to "10% to <20%" policy)

Expected Probability of -20% Cap = (probability assigned to "20% to <30%" policy) + ½(probability assigned to "10% to <20%" policy) + ½(probability assigned to "30% to <40%" policy)

Expected Probability of -40% Cap = (probability assigned to "40% to <50%" policy) + ½(probability assigned to "30% to <40%" policy) + (probability assigned to "50% or greater")

(This page is intentionally left blank)

190

Appendix C: Additional Sensitivity of Backstop Growth Rate Limits
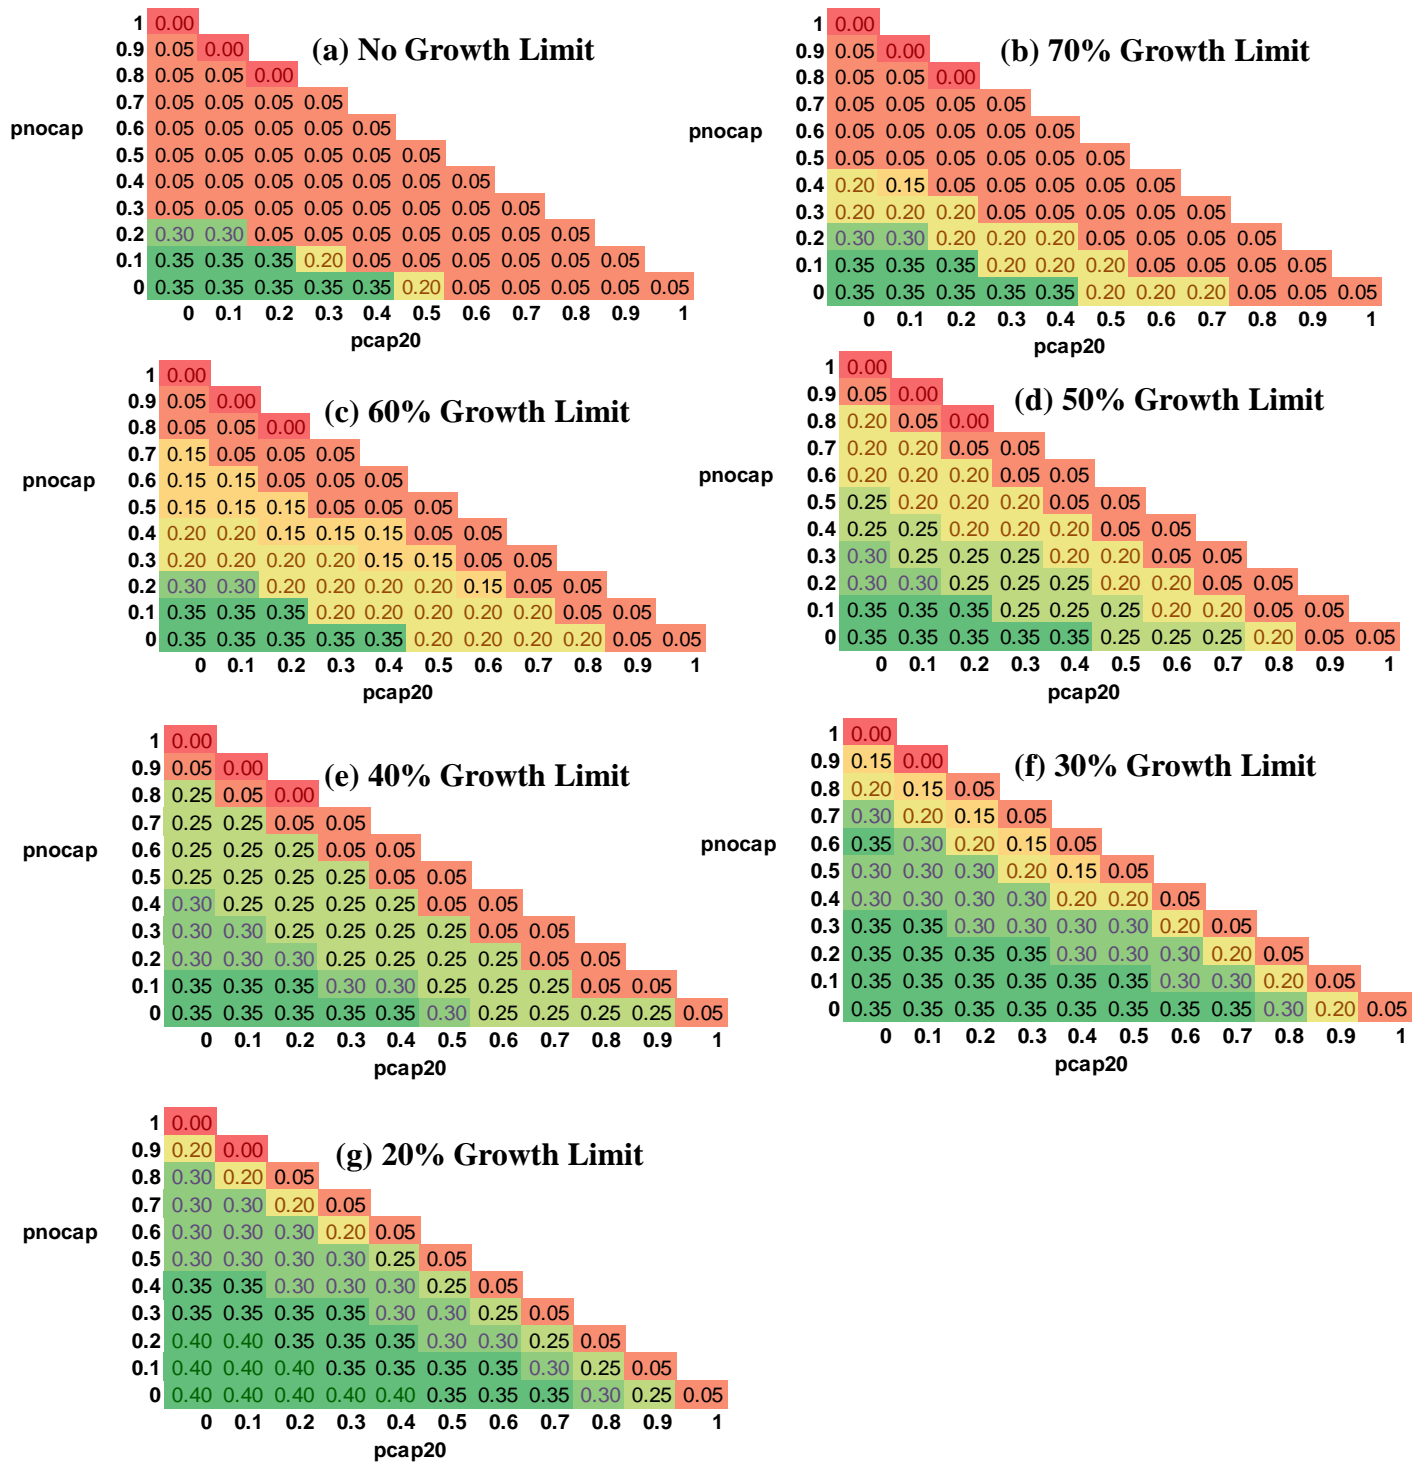
Figure C.1 Optimal Stage 1 Low-Carbon Technology Share by Policy Probability under Different Low-Carbon Growth Rates