

Hedging Strategies: Electricity Investment Decisions under Policy Uncertainty

Jennifer F. Morris,* Vivek Srikrishnan,** Mort D. Webster,** and John M. Reilly***

ABSTRACT

Given uncertainty in long-term carbon reduction goals, how much non-carbon generation should be developed in the near-term? This research investigates the optimal balance between the risk of overinvesting in non-carbon sources that are ultimately not needed and the risk of underinvesting in non-carbon sources and subsequently needing to reduce carbon emissions dramatically. We employ a novel framework that incorporates a computable general equilibrium (CGE) model of the U.S. into a two-stage stochastic approximate dynamic program (ADP) focused on decisions in the electric power sector. We solve the model using an ADP algorithm that is computationally tractable while exploring the decisions and sampling the uncertain carbon limits from continuous distributions. The results of the model demonstrate that an optimal hedge is in the direction of more non-carbon investment in the near-term, in the range of 20–30% of new generation. We also demonstrate that the optimal share of non-carbon generation is increasing in the variance of the uncertainty about the long-term carbon targets, and that with greater uncertainty in the future policy regime, a balanced portfolio of non-carbon, natural gas, and coal generation is desirable.

Keywords: Energy policy, Uncertainty, Electricity, Investment, Emissions, General equilibrium, Approximate dynamic programming

<https://doi.org/10.5547/01956574.39.1.jmor>

1. INTRODUCTION

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 a critical component of any emissions reduction efforts. In the U.S., electric power generation is responsible for approximately 40% of all carbon dioxide (CO₂) 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. Electricity generation investments are expected to operate for 40 or more years, so the decisions we make today can have long-term impacts on the electricity system and the ability to meet long-term environmental goals. Uncertainty in future government climate policy affects the solvency of long-lived capacity investments. If a climate policy is implemented

* Corresponding author. Joint Program on the Science and Policy of Global Change, Massachusetts Institute of Technology, Cambridge, MA, USA. Email: holak@mit.edu.

** John and Willie Leone Family Department of Energy and Mineral Engineering, Pennsylvania State University, State College, PA, USA.

*** Joint Program on the Science and Policy of Global Change, Massachusetts Institute of Technology, Cambridge, MA, USA.

during the lifetime of a power plant, it would greatly affect the cost-effectiveness of that plant, which in turn impacts the value of the investment. Currently, many non-carbon 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. In the absence of current regulations, investment now in these non-carbon technologies are only rational if we consider the prospect of future emissions limits.

Regardless of future greenhouse gas emissions restrictions, electric sector investments continue to be made in the present. Our question is: given the uncertainty in future emissions limits, what level of investments (if any) in the near term should be made into expensive non-carbon generation? Investing now in carbon-free generation technologies has higher near-term costs, but future emissions reductions may be achieved at lower costs if necessary. However, there is also the risk of sunk costs if future reductions are not required.

In the liberalized markets in parts of the U.S. and in other countries, and even in regulated vertically integrated utilities, investment decisions are made by many different agents. Nevertheless, in the absence of long-term economy-wide climate policy, government regulations often attempt to guide the technology mix that results from the aggregate of these decentralized decisions through instruments such as renewable portfolio standards and emissions performance standards that promote or de facto rule out specific technologies (e.g., coal without carbon capture). Our question is focused on the aggregate social welfare maximizing level of near-term investment in high-cost carbon-free technologies under uncertainty in future emissions limits. Exploring this question requires applying uncertainty analysis to an economy-wide model that appropriately represents the welfare impacts on the representative consumer.

There is an extensive literature on mathematical programming-based (e.g., linear programming, mixed integer programming, etc.) generation and/or transmission capacity expansion models applied to broad range of research questions. Many of the generation planning analyses focus on, for example, capacity expansion under competitive markets, with uncertainty in wind penetration, fuel prices, carbon allowances, and/or actions by other market agents (e.g., Ehrenmann and Smeers, 2011; Wogrin et al., 2011). Other studies have focused on generation expansion that incorporates operational flexibility, addressing uncertainty in processes on shorter time scales such as wind availability and demand (e.g., Jin et al., 2014; De Jonghe et al., 2012; Ma et al., 2013; Shortt et al., 2013). Game theoretic and stochastic modeling are combined to address random outages and load forecast errors using Monte Carlo simulation in Roh et al. (2009), and random load and price fluctuations using dynamic programming in Barforoushi et al. (2010). Closer to the question explored here, Reinelt and Keith (2007) develop a stochastic dynamic programming model of firm investment decisions that minimizes the expected present value of future power generation costs under uncertain natural gas and carbon prices, exploring the implications of regulatory uncertainty on generation technology choice and the optimal timing of investment. Park and Baldick (2015) employ a two-stage stochastic programming version of capacity expansion, where the load and demand are the random variables, and explore the impacts of carbon taxes and renewable portfolio standards (RPS). Studies by Mejia-Giraldo and McCalley (2014a, 2014b) use an adjustable robust optimization approach to design a flexible expansion plan with respect to an uncertainty set that includes carbon emissions limit. Kettunen et al. (2011) explore investments in power facilities using a multistage stochastic optimization model with exogenous uncertainty in carbon price. Chronopoulos et al. (2016) use a real options framework to investigate how investment timing and capacity sizing decisions are affected by price and policy uncertainty. Bistline and Weyant (2013) use a stochastic formulation of the MARKAL model to study the optimal generation mix with uncertainty

in which of three scenarios of carbon limits will be realized. Similarly, in the transmission planning context, Munoz et al (2014) also use a stochastic formulation to find the optimal transmission additions under uncertainty in future carbon limits and RPS requirements represented by three scenarios. In general, few capacity expansion or investment planning studies consider uncertainty in future carbon emissions limits, and of those that do, most use only a few scenarios. In addition, all of these studies only examine the electric power sector, not the full economy, and therefore are not able to quantify social welfare impacts

Other studies do use models with a representation of the full economy, which allows substitutions among inputs in other sectors as relative prices change. Of these, a number of studies have used two-stage or multi-stage decision models to explore the optimal near-term climate policy in the context of uncertainty in other factors such as climate damage or abatement costs (e.g. Webster et al., 2012; Lemoine & Traeger, 2011; Gerst et al., 2010; Crost & Traeger, 2010; Webster, 2008a; Webster et al., 2008b; Yohe et al., 2004; Webster, 2002; Kelly & Kolstad, 1999; Nordhaus & Popp, 1997; Kolstad, 1996; Manne and Richels, 1995; Hammitt et al., 1992). However, none of these studies have addressed the question faced by regulators and industry in the present political context: before we know what level of emissions-intensity we will want in the future, what level of non-carbon energy production is socially optimal in the near-term?

We contribute to the literature by using a model that represents the full economy embedded within a stochastic framework where the uncertainty in future carbon emissions limits are sampled from a continuous probability distribution. To implement this, we apply approximate dynamic programming (ADP) to a CGE model to investigate near-term electricity investment decisions under uncertainty in future emissions limits. We first incorporate the CGE model into a two-stage dynamic programming (DP) framework. We then develop and apply an ADP algorithm to solve the model at high resolution to provide insight into non-carbon electricity investments under policy uncertainty.

By using a simple and transparent CGE model that nevertheless represents the essential structure of the economy, we demonstrate the importance of representing 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. Using a CGE model has several advantages: all sectors of the economy are included, important feedbacks such as those from substitution and leakage are captured, prices (e.g., fuel prices) are endogenous, “stranded capital” can be represented, and the net impact on consumers (i.e., social welfare) is estimated. The drawbacks of CGE models are their framing as either myopic (each model period is solved independently) or forward-looking (assuming perfect foresight across all model periods), and the need to aggregate and simplify representations of sectors and technologies. However, incorporating the CGE model into a two-stage approximate dynamic program creates a more realistic temporal structure, and much insight can be gained even with simplified technology representation.

This research demonstrates the value of considering uncertainty—illustrating how uncertainty affects investment strategies as well as the expected cost of future policy. 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 results obtained from this model can provide insight and information about socially optimal near-term electricity investment strategies that will hedge against the risks associated with uncertainty in future policies.

The paper is organized as follows: first, we describe the CGE model (Section 2). We then describe the dynamic programming (DP) approach (Section 3). Section 4 then introduces the ap-

proximate dynamic programming (ADP) approach we developed to overcome the curse of dimensionality that plagues DP, and presents the results from the ADP. Section 5 gives a concluding discussion.

2. THE CGE MODEL

The objective of the CGE model presented here is to provide a simple representation of the U.S. economy while preserving essential features: (1) an electric power sector that can use coal, natural gas and non-carbon emitting sources of generation, and (2) the rest of the economy which uses electricity, other fuels and other inputs to produce final goods for the representative consumer. The reason to keep the model otherwise simplified is because the solution of the sequential decision under uncertainty requires many simulations (on the order of tens of thousands).

Toward this goal, we develop a single region computable general equilibrium (CGE) model approximating the U.S. in terms of overall size and composition of the economy that highlights choices between fossil and non-fossil electricity generation investment decisions. 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, services, etc., comprises the vast majority of the economy. The factors of production included are capital, labor and natural resources (crude oil, coal and natural gas). The base CGE model follows the structure of the MIT Economic Projection and Policy Analysis model (Chen et al., 2016; Paltsev et al., 2005), although considerably simplified.¹

The underlying social accounting matrix (SAM) data is based on GTAP 5 (Hertel, 1997; Dimaranan and McDougall, 2002) data recalibrated to approximate 2010, which is used as the base year for the model. The model is written in the General Algebraic Modeling System (GAMS) programming language and is formulated in MPSGE (Rutherford, 1999). Carbon dioxide (CO₂) emissions are associated with fossil fuel consumption in production and final demand.

Production and consumption functions are represented by nested Constant Elasticity of Substitution (CES) functions.² Production functions for each sector describe the ways in which capital, labor, natural resources and intermediate inputs from other sectors can be used to produce output, and represent the underlying technology through substitution possibilities between the inputs. The consumer utility function describes the preference for each good and service and how they contribute to utility (welfare). The change in aggregate consumption is an equivalent variation measure of welfare in each period.

Two electricity generation technologies are represented: conventional (i.e., fossil) and non-carbon. The single conventional electric technology uses coal and natural gas as its fuel inputs, allowing substitution between the two fuel sources.³ Unlike linear programming-based capacity expansion models, the CGE model represents the flow of economic value over longer periods of time (e.g., one year) parameterized with CES production functions. This production function in-

1. For a detailed description of the model see Morris (2013).

2. For a description of how nested CES functions are used in a CGE model see Chen et al. 2015 and Rausch et al., 2011.

3. Due to the aggregation in the GTAP dataset, conventional electricity aggregates all generation in the base year, including nuclear, hydro and other generation. Nuclear and hydro in the base year can be manually separated from conventional fossil generation. However, we have not done that here, since the goal is a simple representation of generation sources (coal, natural gas and non-carbon). Instead, we assume that all generation in the base year is coal or gas, and consider new investments in nuclear and hydro to fall into the category of non-carbon generation.

cludes an elasticity of substitution between coal and gas inputs, so that the relative ratios can adjust when the relative fuel prices change. We assume an elasticity of substitution of 2.5, which is consistent with many econometric estimates (see Stern, 2012). The non-carbon electricity generation technology produces no carbon emissions and is more expensive, representing advanced non-carbon technologies like wind, solar, carbon capture and storage (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 non-carbon technology is modeled as a single generic technology, highlighting the importance of the relative costs of conventional and non-carbon technologies. The electricity produced from the generalized non-carbon technology is a perfect substitute for conventional electricity. It has a higher cost than conventional generation in the first model period, and its price then evolves endogenously over time. The incremental cost is parameterized in the model by a markup factor, which is the cost relative to the conventional generation against which it competes in the base year. We calculate the levelized cost of electricity (LCOE) for several technologies based on EIA (2015) data: conventional coal is 6.9 cents/kWh, advanced nuclear is 12.8 cents/kWh, natural gas with CCS is 9.5 cents/kWh, wind is 8.7 cents/kWh, and solar PV is 17.5 cents/kWh. Each LCOE is compared to that of conventional coal, resulting in markups of 1.86, 1.38, 1.26, and 2.51 for advanced nuclear, gas with CCS, wind and solar PV, respectively. We use 1.5 as a representative markup for the generalized non-carbon technology, meaning the non-carbon technology is assumed to be 50% more expensive than conventional electricity.

The CGE model is dynamic, running from 2010 to 2030 in 5-year time steps. 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 non-carbon electricity technology, (4) population growth, and (5) energy efficiency improvements. The first three processes are endogenous while the last two are exogenous.

Of particular importance for this analysis is the representation of capital vintaging, which is applied to the electricity sector and reflects the irreversibility of decisions. 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 characteristics of that capacity (e.g., amount capital vs. labor vs. fuel, etc.). As the model steps forward in time it preserves four vintages (20 years) of rigid capital, minus depreciation, with each vintage retaining the 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). This approach captures the phenomenon that today’s decisions about how much of each technology to build will affect the electricity system long into the future.

The availability of the non-carbon technology is also an important dynamic. As noted by Jacoby et al. (2006), penetration rates for new technologies have historically exhibited 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 approximate this dynamic, a fixed factor resource is included in the model, which simulates 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 new capacity is installed, the fixed factor resource grows as a function of the technology’s output in the previous period, simulating the effect that increasing production enables that technology to expand further. As non-carbon electricity production expands over time, the fixed factor endowment is increased, and it becomes non-binding. 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 $t + 1$ (for a more detailed explanation see Morris et al., 2014).

3. DYNAMIC PROGRAMMING APPROACH

3.1 DP-CGE Model

The problem we wish to consider is how uncertainty in the future emissions limit impacts the near-term electricity technology mix, as well as near-term emissions decisions. This problem is one of decision making under uncertainty, and requires a stochastic formulation. There are, of course, several approaches to stochastic optimization with recourse. Because the underlying subproblem (the CGE model) is not a linear program, implementation as a stochastic dynamic program has advantages over stochastic programming for this application. In this section, we formulate the problem as a two-stage finite horizon stochastic dynamic programming (DP) problem, with uncertainty about the future emissions limit that is resolved in the second period.

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 endogenously chooses an electricity technology mix (and all other outputs) based on the current-period maximization of consumption. However, here we are interested in the technology mix in each period that maximizes the current period consumption *plus* the *expected future* consumption. Specifically, we solve the following maximization problem:

$$\max_{x_1} C_1(x_1) + E_{\theta} \{ \max_{x_2} [C_2(x_2, S_2, \theta)] \} \quad (1)$$

where C_t is the consumption in period t , x_t is the vector of decision variables, θ is the emissions cap (uncertain in period 1 and becomes known in period 2), S_2 is the system state, which is a function of the previous period's decision, and $t \in \{1, 2\}$.

We initially solve the maximization problem in equation (1) using dynamic programming (DP). The DP objective is to choose actions to maximize total expected discounted social welfare in the economy over the planning horizon. The Bellman equation (Bellman, 1957) expresses the optimality conditions as:

$$V_t = \max_{x_t} [C_t(S_t, x_t) + \gamma E \{ V_{t+1}(S_{t+1}(S_t, x_t, \theta)) \}] \quad (2)$$

where t is the decision stage, V is the total value, S is the vector of state variables (electric power capacity of each technology and cumulative emissions level), C is the economy-wide consumption (welfare), x is the vector of decision variables (non-carbon share of new electricity and amount of emissions reductions), θ is the uncertainty (probability distribution over Stage 2 carbon constraint), and γ is the discount factor = $(1 - \text{discount rate})$, where the discount rate = 4%.

In the DP, there is uncertainty (θ) in the future climate policy (i.e., the cumulative emissions limit in Stage 2). Two decisions are made: (1) the non-carbon technology's share of *new* electricity in each stage (i.e. how much of the new electricity capacity built should consist of non-carbon technologies?), and (2) Stage 1 reductions of electricity emissions (i.e. is it cost-effective to begin reducing emissions now in anticipation of future policy?). The first decision about the non-carbon production share in each stage is exogenously imposed as a constraint within the CGE model. This

constraint forces the allocation of new capital and of labor to the non-carbon sector in amounts that may not have been optimal in a static myopic CGE model. The second decision to reduce Stage 1 emissions via a “self-imposed” emissions cap provides a price signal (i.e. the carbon price) in the CGE model that affects the operation of existing electricity capacity as well as the optimal share of non-carbon new electricity, by shifting production among sectors, substituting input fuels (gas for coal), and/or reducing electricity output. One potential result is to leave some vintage capital unused or underutilized. Ultimately, this emissions reduction decision variable affects the choices of coal vs. natural gas, conventional vs. non-carbon, and building new vs. operating existing capacity differently, as well as total amount of electricity output.

The uncertainty modeled is the level of the Stage 2 emissions cap. We define emissions caps in terms of the *cumulative* emissions from the *electric power sector* from 2015 to 2030 (e.g. an emissions cap of X% below the cumulative no policy emissions). The policy cases focus on cumulative emissions because it is cumulative emissions, not the specific emissions level at any one point in time, which determine long-term global temperatures. This framing is also consistent with previous cap-and-trade proposals in the U.S., which have included intertemporal flexibility (banking and borrowing of emission permits) for meeting a long-term emissions target.

Because this uncertainty aggregates the future results of many political, social, and scientific processes, there is no objective source for a probability distribution over future emissions caps. To provide insights into the effects of uncertainty on near-term decisions, we explore several illustrative probability distributions and investigate the relative effect of alternative beliefs on the near-term optimal decision in online Appendix A. Using Approximate Dynamic Programming in Section 4, we treat this uncertainty as continuous, and compare the impacts of different mean-preserving spreads.

The underlying CGE model consists of 5-year time steps, but the time horizon for the DP aggregates these into two decision stages. Stage 1 includes CGE periods 2015 and 2020 while Stage 2 includes 2025 and 2030 (2010 is the benchmark year). 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 firms maximizing their own profits according to their production functions, input costs and the policy constraints imposed by the central planner under conditions of perfect competition. The decisions and uncertain outcomes are exogenously imposed on the CGE model, which then endogenously chooses all other output quantities, including the shares of natural gas and coal generation. Consumption (social welfare) is calculated for each decision stage and used in the Bellman equation (eq. 2), along with the probabilities of uncertain outcomes, to identify optimal Stage 1 decisions. The model formulation follows the classic act-then-learn framework (Manne and Richels, 1992): Stage 1 decisions are made under uncertainty in the emissions cap, which is revealed before the Stage 2 decision is made. In effect, the CGE model performs intra-period optimization and the DP performs inter-period optimization.

3.2 Representing Path Dependency

In the DP formulation of the electricity investment decision problem, we include an additional element to the problem structure—that of path-dependency in the non-carbon investment decision. For many of the non-carbon technologies, scaling up to a large penetration of the new technology within the energy sector requires a significant investment in new physical/organizational infrastructure. For renewables such as wind and solar, their intermittency will require fundamental restructuring of power systems operations/markets to maintain reliability. For carbon capture and

sequestration, a new infrastructure for transportation and storage of carbon must develop. For such technologies, insufficient levels of near-term investment would likely constrain the rate at which the share of the technologies can be expanded in the future, even if needed.

We represent this effect with a parameter defined as the maximum allowable increase in the share of non-carbon production in Stage 2, relative to the share of production in Stage 1. This exogenous non-carbon production growth rate limit may be binding even when the fixed factor is no longer binding, depending on the assumed maximum growth rate. This additional constraint limits the rate of growth of non-carbon production as a share of new electricity between Stage 1 and Stage 2. For example, if the share of non-carbon cannot increase by more than 50 percentage points from Stage 1 to Stage 2, and if the production share was 0% in Stage 1, then 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 non-carbon grows—non-carbon could constitute 0% of new electricity in Stage 1 and 100% in Stage 2. This would represent a situation in which all new electricity production added during 2021–2030 is non-carbon. While theoretically possible, such a solution does not seem likely or technologically feasible. All investors would have to decide to build non-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 non-carbon capacity additions. However, in the past we have seen fairly rapid expansion of nuclear electricity, and currently natural gas generation is rapidly expanding due to the new supply of shale gas, which suggests that there may not be a limit to the rate of non-carbon electricity growth. Because it is difficult to assess and people have widely varying opinions about what type of non-carbon growth rate is realistic from engineering and technological perspectives, we perform sensitivity analysis to explore a range of assumptions about the limit to non-carbon generation expansion rates.

4. APPROXIMATE DYNAMIC PROGRAMMING APPROACH

The focus of this analysis is on the optimal near-term decisions about new electricity investment shares and emissions reductions under different scenarios of uncertainty in the future emissions limit. As a first step, we solved the exact DP, which required discretizing the decision and uncertainty spaces (results from the exact DP are presented in online Appendix A). Computational demands cause the implementation to be relatively coarse resolution. To gain better insight into the true decision problem, it is desirable to represent the decision and uncertainty spaces as continuous.

Here, we apply Approximate Dynamic Programming (ADP),⁴ which is a class of methods developed to address the Curse of Dimensionality that plagues dynamic programming (Powell, 2011; Bertsekas & Tsitsiklis, 1996). ADP allows for higher resolution and larger-scale representation of problems, while remaining computationally feasible. ADP combines Monte Carlo simulation, adaptive sampling, and response surface approximation strategies to approximate the value function (eq. 2). The idea is that instead of exhaustively searching through all possible states, decisions, and information signals, ADP iteratively samples paths through the scenario tree to construct an approximation of the value function, which can then be used to make optimal decisions

4. ADP is also sometimes referred to as neuro-dynamic programming, forward dynamic programming and adaptive dynamic programming.

for any feasible state. In turn, those approximate decisions guide the next set of sample paths to progressively focus on the highest value region of the state space. The optimization is broken into two steps: (1) deterministic optimization to select the best action for a sample state and sample information signal based on the current approximation of the value function, and (2) stochastic simulation to improve the estimate of the expected value for a sample state, given its best action. Recent work (e.g., Webster et al., 2012a & 2017; Godfrey & Powell, 2002; Basler, 2006; Powell et al., 2012) has successfully implemented ADP methods on large-scale, multi-dimensional problems.

There are two broad classes of ADP algorithms: (1) iterative approaches, which are based on value iteration or policy iteration (examples include TD-learning and Q-learning), and (2) linear programming based approaches (Schweitzer and Seidman, 1985; de Farias and Van Roy, 2003). Here, we develop and apply an ADP solution method that combines the strengths of both classes of algorithms. We call our algorithm a *Q*-factor Adaptive Relaxed Linear Problem (QARLP) algorithm. A full description of the algorithm is presented in online Appendix B. We apply the QARLP algorithm to solve the stochastic dynamic program formulated in the previous section for continuous decision and state spaces.


4.1 ADP Results

Using our QARLP algorithm to explore our problem at a higher resolution, we provide insight into the effect of variance in the uncertainty of future climate policy, the impact of constraints on non-carbon expansion, how the optimal strategy compares to common heuristics, and the value of including uncertainty.

4.1.1 Role of Variance in Uncertainty

We first explore how the optimal Stage 1 investment is impacted by the variance in the uncertainty of the Stage 2 emissions limit. To model the continuous distribution of future emissions limits, we use the Beta distribution family. The Beta distribution is a flexible parametric distribution with two parameters, α and β , and is defined over the closed interval $[0,1]$, but can be scaled and shifted to any other interval. We explore beta distributions with shape parameters $\alpha = \beta = 0.25, 0.5, 1, 1.5, 2,$ and 4 , and all are location-scaled onto the interval $[0.6; 1]$ (where $0.6 = -40\%$ cap and $1 = \text{no cap}$) to allow for a maximum emissions cap of -40% from the reference case. The probability density functions for these distributions are shown in Figure 1. These distributions all have the same mean of 0.8 (-20% cap), but differing variances. Also note that the Beta(1,1) distribution is by definition equivalent to the Uniform(0,1) distribution.

The optimal decisions with respect to each of these distributions are given in Table 1. Increasing the variance in the uncertainty of the Stage 2 emissions limit leads to a greater share of non-carbon investment and to more emissions reductions being optimal in Stage 1. Under the highest variance distribution, the optimal decision is 29% non-carbon electricity production from new capacity and 16.1% emissions reductions below the reference emissions. Under the lowest variance distribution, a 5% share of non-carbon electricity and 12.3% emissions reductions is optimal. The more aggressive reductions under the high variance distribution is a result of the higher probability that the Stage 2 cap will be stringent. Because the marginal costs of emissions reductions are increasing, there is an asymmetric loss function such that it is very costly to meet a strict cap in Stage 2 if little or no non-carbon investment or emissions reductions have occurred in Stage 1. In order to avoid that situation, more near-term investment in non-carbon generation and emissions

Figure 1: Distributions of Stage 2 Emissions Limit with Constant Mean and Different Variances**Table 1: Optimal Stage 1 Emissions and New Electricity Investments for Different Distributions of the Stage 2 Emissions Limit**

Distribution	Variance	Emissions Reductions	Low-Carbon Share	Coal Share	Natural Gas Share
Beta(0.25,0.25)	0.167	16.1%	29%	24%	47%
Beta(0.5,0.5)	0.125	15.3%	25%	24%	51%
Beta(1,1)	0.083	15.0%	20%	22%	58%
Beta(2,2)	0.050	13.5%	5%	19%	76%
Beta(4,4)	0.028	12.3%	5%	20%	75%

Note: Non-carbon growth limit is 0.3.

reductions is optimal as the probability of a strict cap increases. As a consequence, the share of natural gas decreases with the variance of the Stage 2 distribution while the share of coal increases (though to a lesser extent) with the variance. These results demonstrate that the main tradeoff is between non-carbon and natural gas generation, and also that an increasingly well-balanced portfolio approach is optimal as the variance of the Stage 2 distribution increases. Under high variance, high non-carbon investment provides a hedge against a strict cap while higher coal investment provides a balancing hedge against a lax or no cap, and natural gas fills in the rest. As the variance decreases and a moderate cap becomes more likely, the investment portfolio increasingly relies on natural gas.

4.1.2 Role of Constraints on Non-Carbon Expansion

Here we demonstrate how constraints on the ability to expand the non-carbon electricity share in the future affect the optimal Stage 1 decisions. As described above, past experience with

Table 2: Sensitivity of Optimal Stage 1 Decisions to the Non-carbon Growth Limit

Distribution	Growth Limit 0.3		Growth Limit 0.5		Growth Limit 0.7	
	Non-Carbon Share	Emissions Reduction	Non-Carbon Share	Emissions Reduction	Non-Carbon Share	Emissions Reduction
Beta(0.25,0.25)	29%	16.1%	21%	16.0%	5%	15.0%
Beta(0.5,0.5)	25%	15.3%	18%	15.2%	5%	14.1%
Beta(1,1)	20%	15.0%	5%	13.8%	5%	12.5%
Beta(2,2)	5%	13.5%	5%	12.3%	5%	10.1%
Beta(4,4)	5%	12.3%	5%	10.2%	5%	7.5%

technological transitions suggests that technologies with a low share in Stage 1 may not be able to rapidly expand to a larger share of new generation in Stage 2, due to infrastructure limits, capacity constraints, and electricity system adjustment costs. We have modeled this non-carbon growth limit as a constraint on the maximum allowable increase in the renewable share of electricity from Stage 1 to Stage 2. This growth limit is an additional constraint to the fixed factor constraint that is endogenous in the CGE model. We compare the results obtained for growth limits of 0.3, 0.5 and 0.7 (a growth limit of 0.5 means, for example, that if 20% of investment in Stage 1 is in the non-carbon technology, then the max non-carbon share in Stage 2 is 70%).

Table 2 gives the resulting optimal Stage 1 decisions for various non-carbon growth limits and distributions of Stage 2 emissions limit with different variances. The results for the 0.3 growth limit are the same as those presented above in Table 1. With a higher growth limit—more flexibility to expand the non-carbon share in the future despite minimal investment in the near-term—a lower share of non-carbon electricity and fewer emissions reductions are optimal in Stage 1. This result is consistent across all distributions of uncertainty in future emissions limits. In this model, a 0.7 growth limit is not a binding constraint in any scenario. When non-carbon growth is unconstrained, the optimal Stage 1 non-carbon investment share is 5% for all distributions. 5% is the minimum non-carbon production share needed to render the fixed factor constraint in the CGE model non-binding in Stage 2, thereby providing unlimited flexibility to expand non-carbon penetration to the desired amount in Stage 2. Greater flexibility to expand non-carbon in Stage 2 also requires fewer emissions reductions in Stage 1, because emissions reductions can be achieved more easily in stage 2 by increasing the share of non-carbon generation.

4.1.3 Optimal vs. Heuristic Strategies (Consequences of Alternative Rules)

An important question is how the optimal strategy identified by a decision-making under uncertainty approach compares with common heuristic strategies. To address this question we compare the results from four different strategies, each with different optimal Stage 1 decisions:

- 1) “Optimal”: Optimal decision under uncertainty found using QARLP
- 2) “Deterministic”: Heuristic strategy that assumes that the expected value of the Stage 2 emissions limit distribution will be the certain outcome (i.e., the Flaw of Averages)
- 3) “Myopic”: Heuristic strategy focused on maximizing near-term welfare, such that Stage 1 involves no emissions reductions and the minimal non-carbon share needed to overcome the fixed factor constraint
- 4) “Aggressive”: Heuristic strategy based on the precautionary principle to avoid the worst case outcome in the future (i.e., a strict cap with little to no prior non-carbon or emissions reductions investments made).

Table 3: Stage 1 Decisions under Different Heuristic Strategies

Strategy	Stage 1			
	Emissions Reduction	Non-Carbon Share	Coal Share	Natural Gas Share
Optimal	15.03%	20%	22%	58%
Deterministic	7.60%	5%	28%	67%
Myopic	0.01%	5%	39%	56%
Aggressive	24.24%	50%	21%	29%

Note: These Optimal strategy results assume a Beta(1,1) distribution for the Stage 2 emissions limit and a non-carbon growth limit of 0.3.


Table 4: Stage 2 Non-carbon Share given Different Stage 1 Heuristic Strategies as a Function of the Revealed Stage 2 Emissions Limit and the Non-carbon Growth Limit

Revealed Emissions Limit	Non-Carbon Growth Limit = 0.3				Non-Carbon Growth Limit = 0.5			
	Optimal	Deterministic	Myopic	Aggressive	Optimal	Deterministic	Myopic	Aggressive
0.6	0.5	0.35	0.35	0.54	0.55	0.55	0.55	0.54
0.7	0.49	0.35	0.35	0.34	0.55	0.55	0.55	0.34
0.8	0.34	0.35	0.35	0.03	0.38	0.4	0.54	0.03
0.9	0.03	0.22	0.29	0.02	0.11	0.22	0.29	0.02
1	0.03	0.03	0.04	0.02	0.03	0.03	0.04	0.02

Each of these strategies corresponds to different objectives. The Optimal strategy, the focus of this analysis, derives from the solution to the approximate dynamic program that explicitly seeks to maximize expected welfare given the uncertainty in the Stage 2 emissions cap and the ability to adapt in Stage 2. The Deterministic (expected value) strategy is often used when it is deemed that incorporating a probability distribution is computationally prohibitive or would be too difficult to incorporate into the model. The Myopic strategy focuses on short-term welfare, and does not consider the potential future costs. The Aggressive strategy takes a precautionary approach in which the goal is to minimize the worst-case scenario (which in our case is the most stringent cap in Stage 2, assumed to occur with certainty). Note that determining the decisions under each strategy (other than Optimal) is straightforward and requires much less computational effort than the DP/ADP solutions.

The optimal Stage 1 decisions from following each strategy are given in Table 3, and the respective Stage 2 non-carbon shares are given in Table 4 for several possible revealed emissions limits. In some cases, the non-carbon growth limit has a clear impact and is a binding constraint. For example, the Deterministic and Myopic strategies are extremely limited in their Stage 2 non-carbon share expansion under a 0.3 growth limit, and are moderately constrained even under a 0.5 growth limit. Even if the Optimal decision is made in Stage 1, Stage 2 non-carbon expansion can still be limited by the growth limit constraint, depending on the realized cap. The range of the caps for which the Stage 2 policies are constrained by the growth limit is a function of not only the Stage 1 non-carbon share, but also the Stage 1 emissions reductions, as seen by the differences in the Deterministic and Myopic curves. Under the Myopic strategy, the higher emissions in Stage 1 requires all of the cumulative emissions reductions to occur in Stage 2. Further, the inability to increase the non-carbon share in Stage 2, because of insufficient investment in Stage 1, leads to

Figure 2: Cumulative Consumption Given Different Stage 1 Strategies and Depending on the Revealed Stage 2 Cap: (a) 0.3 Non-carbon Growth Limit, and (b) 0.5 Non-carbon Growth Limit




even higher marginal costs of those emissions reductions, which must therefore be achieved from conventional electricity generation.

For each of the four strategies, we perform a Monte Carlo simulation, sampling the Stage 2 emissions cap to obtain probability distributions of costs and other outcomes. The results here assume the uncertainty in the cumulative emissions cap is uniform (i.e., $\text{Beta}(1,1)$), and we perform Latin Hypercube Sampling with sample sizes of 1000. The full impacts of a regulation on the electric sector cannot be estimated by focusing only on that sector. Those regulations will shift the quantities demanded and produced, which in turn causes a shift in prices. The economy will adjust to the new prices through substitution among inputs to other production sectors and through substitution by the consumer across the consumption bundle. The ultimate impact on consumption can be amplified or dampened by the substitution effects. Moreover, the net change in carbon emissions may be less than the electric sector target due to leakage to other sectors. The CGE model used in this analysis is intended to capture all of these effects.

Figure 2 shows the resulting percentage change in cumulative (2010–2030) consumption (social welfare) relative to the reference no policy consumption, using a discount rate of 4%. The Optimal strategy results in smaller consumption losses than the Aggressive decision for all but the strictest caps. The greater losses in cumulative consumption under the Deterministic and Myopic strategies as the cap increases in stringency demonstrates the asymmetric behavior discussed above. With increased flexibility to increase the renewable share (growth limit of 0.5), these losses are not as large, but the asymmetric trend is the same.

Figure 3 uses the probability distribution of the Stage 2 emissions limit to translate the cost curves from Figure 2 into boxplots. The Optimal strategy results in the lowest expected cost (reduction in consumption) for both non-carbon growth limits. The Optimal strategy also exhibits smaller variance in costs than the Deterministic and Myopic strategies, significantly reducing the downside (i.e., the largest cost is not as large as the other strategies) without limiting the upside by

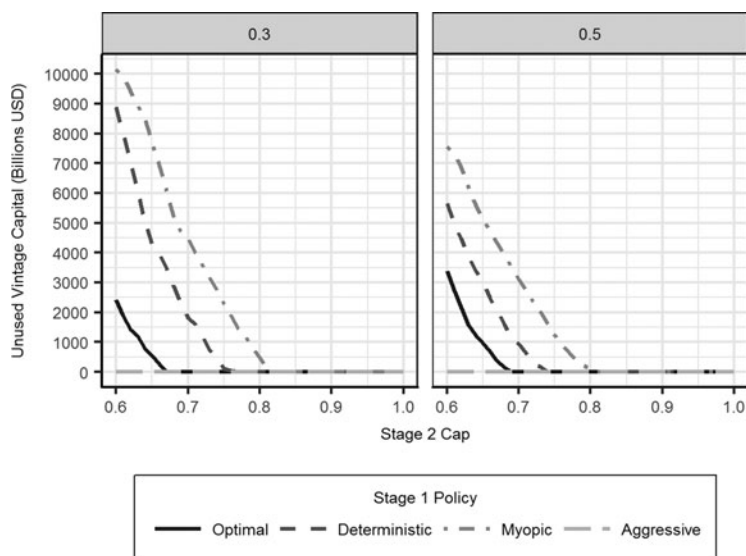
Figure 3: Loss in Cumulative Consumption by Strategy Relative to Reference: (a) 0.3 Non-carbon Growth Limit, and (b) 0.5 Non-carbon Growth Limit



too much. Although the Aggressive strategy has the smallest variance in cost, the smallest worst case loss, and a lower expected cost than the Deterministic and Myopic strategies, it also has the greatest minimum cost (i.e., best case outcome is not as good as the best case outcomes under the other strategies). That is the nature of a precautionary or robust strategy—it protects from the downside but also truncates the upside. It is a very risk-averse strategy that flattens out the possible consumption outcomes (as apparent in Figure 2). The Myopic strategy has the greatest variance in cost, and, with the greatest potential for a poor outcome, has the greatest expected cost. The Deterministic strategy only fares slightly better than the Myopic. For both, the downside is driven by the potential of ending up in the worst case scenario—a strict cap in Stage 2 after little or no Stage 1 non-carbon investments or emissions reductions. The increased flexibility that comes with a high non-carbon growth limit, does not change these qualitative relationships, but does reduce both the variance (particular on the low end) and the expected cost for the three strategies that are impacted by the growth limit (all but the Aggressive strategy).

By examining the utilization of vintaged capital, tracked by the CGE model, the drivers for the asymmetric losses are more apparent. Figure 4 shows unused vintaged capital in Stage 2 from the electric sector (aggregating non-carbon and fossil capital) as a function of the Stage 2 cap and the Stage 1 decision strategy. As the cap becomes more stringent, the Deterministic and Myopic strategies (and the Optimal strategy to a lesser extent) lead to increasing amounts of vintaged capital being stranded. This represents overinvestment in conventional fossil generation capacity in Stage 1 that cannot be fully utilized in Stage 2 due to the emissions constraint. It is very expensive to leave existing (vintaged) capacity unused or underutilized, and this cost contributes to the greater consumption losses for stringent caps evident in Figures 2 and 3. In contrast, even when there turns out to be a less binding emissions cap or no cap, the Aggressive and Optimal strategies continue to utilize all vintaged non-carbon generation. It is never economic to leave the non-carbon electricity capital unused, even if it is not required to meet the emissions limit. This difference between stranded fossil capital under tight caps and the absence of stranded non-carbon capital under loose caps causes the asymmetry in consumption losses.


Figure 4: Unused Vintage Electricity (Fossil + Non-Carbon) Capital for the Stage 1 Heuristic Strategies as a Function of the Revealed Stage 2 Cap: (a) 0.3 Non-carbon Growth Limit, and (b) 0.5 Non-carbon Growth Limit



The results above demonstrate the cost asymmetry between overinvesting in conventional generation and overinvesting in non-carbon generation. This is clearly shown by comparing the Myopic strategy under a strict cap (e.g. 0.6, or -40%) to the Aggressive strategy under no cap (e.g. 1.0). The Myopic strategy with a realized cap of 0.6 in Stage 2 results in a consumption loss of 0.9% relative to the reference (no policy) case, and \$10 trillion worth of vintaged fossil capital stranded. The Aggressive strategy with no cap in Stage 2 results in only 0.2% consumption loss relative to reference and no vintage capital stranded. In that case, even though the non-carbon generation turned out to be unnecessary in the absence of policy, it continues to be used. This asymmetry is a consequence of the variable costs of electricity production. With an emissions limit, the implicit carbon price increases the fuel cost component of conventional (fossil) generation. Even though the capital investment is a sunk cost, the variable cost of operating the conventional generation (mainly the fuel cost) becomes greater than the full cost of investing in new non-carbon generation, and therefore goes unused. On the other hand, non-carbon generation has low variable costs (and no fuel costs), so once the capital investment is made, operation is relatively inexpensive, and is lower cost than the full cost of investing in new conventional generation. As a result, vintage non-carbon capacity continues to be used even when there is no emissions limit. Because of this cost asymmetry a robust near-term investment strategy should err on the side of more non-carbon generation rather than more carbon-emitting generation.

In addition to providing estimates of consumption (social welfare) and stranded capital, our CGE approach also has the advantage of being able to track emissions leakage. Because we have defined our policies as limits to emissions from the electricity sector, emissions in other sectors of the economy that are not covered by the cap can increase relative to what those emissions would have been without any policy. The amount of leakage largely depends on the ability to shift dirtier energy sources to unregulated sectors of the economy. A cap on electricity sector emissions will decrease the amount of fossil fuel energy, particularly coal, used in generation. This may cause the

Figure 5: Emissions Leakage Given Different Stage 1 Strategies and Depending on the Revealed Stage 2 Cap: (a) 0.3 Non-carbon Growth Limit, and (b) 0.5 Non-carbon Growth Limit



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 lower cost fossil resources helps to offset the cost of the sectoral policy, but also reduces the economy-wide emissions impact of the sectoral cap. Figure 5 shows emissions leakage for each of the Stage 1 strategies depending on the Stage 2 cap and the assumed non-carbon growth constraint. In general, the more stringent the policy, the more leakage occurs. However, for stringent caps (less than 0.8), the Myopic and Deterministic strategies lead to greater leakage. This is because the absence of sufficient investment in the first stage necessitates dramatic reductions in Stage 2, which increases the electricity prices significantly. Higher electricity prices induce substitution away from electricity to other sectors that do not have emissions limits, resulting in leakage.

4.2 Value of Including Uncertainty

One of the contributions of this research is demonstrating how decision making under uncertainty can be represented using a CGE model and the value of doing so. The expected value of including uncertainty (EVIU) (Morgan and Henrion, 1990; Savage, 2009), also called the value of a stochastic solution (VSS) (Birge and Louveaux, 1997), is a metric that captures the value of representing uncertainty or, equivalently, the incremental cost of choosing a non-optimal strategy by explicitly neglecting uncertainty. EVIU reflects the improvement in decisions that can be obtained by formally modeling uncertainty in the decision-making process. To quantify the value of accounting for uncertainty, we compare the expected value of the Optimal strategy to that under the other strategies, for three different probability distributions for the Stage 2 cap, and for two different non-carbon growth limits (Table 5). The EVIU is calculated by comparing policy costs from each strategy to those from the Optimal strategy. Note that the definition of EVIU typically

Table 5: Expected Policy Cost and Expected Value of Including Uncertainty (EVIU)

Distribution	Growth Limit	Expected Policy Cost for Each Strategy (Relative to Reference) (\$B)				EVIU (% Increase in Cost of Alternative Strategy vs. Optimal Strategy)		
		Optimal	Deterministic	Myopic	Aggressive	Deterministic	Myopic	Aggressive
Beta(0.5,0.5)	0.3	149	228	312	183	53%	109%	23%
	0.5	137	174	233	183	27%	70%	34%
Beta(1,1)	0.3	134	181	256	179	35%	91%	34%
	0.5	121	143	199	179	18%	64%	48%
Beta(2,2)	0.3	121	143	213	174	18%	76%	44%
	0.5	103	118	171	174	15%	66%	69%

refers to the comparison of the “Deterministic” strategy to the optimal; i.e., the loss from using expected values for the uncertain parameters and solving for the deterministic optimal decision for that case. However, “Myopic” and “Aggressive” represent other commonly used heuristics that are either explicitly or implicitly employed as an alternative to formally modeling the uncertainty.

The expected costs of each strategy (difference between expected cumulative consumption under the strategy and under the reference case in billions of U.S. \$) are given in Table 5 for each strategy, as well as the percentage increase in costs from the three strategies that do not explicitly solve the stochastic decision problem. For the Deterministic strategy, the consumption losses from using this heuristic range from 15% to over 50%, depending on the variance in the uncertainty and the degree to which current non-carbon investment constrains future non-carbon investment. The greatest losses occur for the highest variance distribution—the less confident one is in the direction of future policy, the greater the expected losses from using this analytical shortcut. The losses from the Myopic strategy are even greater, ranging from 60% to more than 100%, and are also greatest for higher variance distributions. Under the Aggressive strategy, the trend is reversed in that the expected losses are greatest for lower variance distributions. This is because the Aggressive strategy is a heuristic for finding a robust strategy; this is visually demonstrated in Figure 2 where the costs of this strategy are nearly constant over the full range of potential Stage 2 emissions limits.

5. CONCLUSIONS

When making a decision under uncertainty, one often needs to hedge against a risk. More difficult are decisions that must balance two risks, each of which imply hedging in different directions. The problem analyzed here is an example of that type of decision problem. The amount of non-carbon electricity generation that should be developed over the next decade, before we know what long-term carbon targets are, is a decision that must balance the risk of overinvesting in non-carbon energy with the risk of underinvesting. Underinvesting could lead to very high costs in the future if dramatic carbon reductions are needed. Overinvesting could impose higher costs now that turn out to be unnecessary. In the absence of careful analysis, it is not obvious a priori in which direction society should be hedging.

The work presented here has applied a simple computable general equilibrium model of the U.S. economy, and embedded it within a stochastic approximate dynamic programming framework. The dynamics of this model have demonstrated that there is in fact an asymmetry to the costs of overinvesting in non-carbon generation vs. underinvesting. Specifically, the expected losses from underinvesting in non-carbon electricity in the near-term is much larger than the corresponding

costs of overinvesting. There are three main factors that drive this asymmetry in costs. First, excessive investment in conventional fossil technologies poses a great risk under policy uncertainty—if a stringent policy is implemented in the future, some of the overbuilt conventional fossil capacity may become stranded due to high operating costs, driven by high fuel costs that reflect the shadow price of the carbon limit. In the converse situation, excessive non-carbon generation capacity built in the near-term will be used in the long term even in the absence of stringent emissions caps. There will be unnecessary sunk capital costs, but low operating costs would allow this non-carbon capacity to continue to be used. Second, insufficient investment in non-carbon generation in the near-term could reduce the flexibility to ramp up investment later, even if needed, due to capacity constraints and a lack of infrastructure. Third, if non-carbon generation capacity was insufficient in the near term and future expansion is therefore constrained, dramatic reductions, if needed rapidly, would have to be achieved by other more costly means. Because of this cost asymmetry a robust near-term investment strategy should err on the side of more non-carbon generation rather than more carbon-emitting generation.

Another result of this analysis is that the optimal share of new non-carbon generation in the next decade is increasing in the variance of the uncertainty in future carbon targets. In general, the less confidence we have in what the long-term carbon reduction goals will turn out to be, the more non-carbon generation we should build now. However, equally critical is that with greater uncertainty, more coal generation should also be added, at the expense of natural gas. New generation dominated by natural gas is optimal if we have high confidence that the long term carbon goal will be a moderate one (e.g., 20% reduction by 2030). Greater uncertainty is best addressed by a balanced portfolio of a range of technologies with varying emissions rates in the near-term.

Overall, while our model is stylized and numbers should be treated with caution, our results show that near-term non-carbon investments should be in the range of 5–30% of new production and emissions reductions in the range of 8–16%, depending on the uncertainty in future emissions policies and constraints on future non-carbon expansion. The low end of the ranges is best if we expect that a future strict cap is quite unlikely, or one in which it is easy to quickly expand non-carbon generation. The high end of the ranges is best if one believes that there is at least a non-negligible probability of a future strict cap, especially if the ability to expand non-carbon generation could depend on what we put in place now.

Finally, this work demonstrates the value of formally representing uncertainty in decision support models. Existing modeling approaches typically do one of two things with regard to uncertainty: (1) ignore it (either by being myopic or pretending we have perfect foresight), or (2) consider scenarios, each with certainty, and then use the optimal decision from the middle scenario or from a scenario that assumes the average value of the uncertainty to be the known value. Here we compared the optimal decisions identified by our approximate dynamic programming approach to other common strategies, including a myopic “wait and see” approach, a deterministic expected value or “flaw of averages” approach, and an aggressive precautionary or “worst case” approach. We showed that these approaches result in different near-term investment strategies that increase the expected policy cost relative to the optimal hedging strategy from an approximate dynamic programming framework that explicitly considers decision making under uncertainty. Given the experimental design, the expected policy cost is increased by 15–54% by pursuing a deterministic strategy, by 64–109% by pursuing a myopic strategy, and by 23–69% by pursuing a worst-case strategy. Considering uncertainty in decision-making results in investment strategies that minimize expected policy costs, and are reasonably robust while not being as extreme as more precautionary approaches.

The main technical contribution of this work 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 (Monte Carlo), but without capturing the critical nature of 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. It does so by first incorporating a CGE model into a dynamic programming framework, and then applying an algorithm for an approximate dynamic programming version of the model that can replicate the DP results with accelerated computation time. Our *Q*-Factor Adaptive Relaxed Linear Programming (QARLP) approach also allows the exploration of the model at dramatically higher resolution than the DP, permitting greater insight into optimal near-term decisions. In addition, this approach demonstrates that the model can be scaled up without being limited by the curse of dimensionality. Future work using this approach could use a more detailed/disaggregated model structure, more technologies, or more decision periods. Ultimately, applying this new modeling approach allows for a quantitative evaluation of near-term electric power investment decisions under uncertainty. In doing so, we demonstrate how a CGE model can be structured 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.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the financial support for this work provided by the U.S. Department of Energy, Office of Science under grants DE-SC0003906 and DE-FG02-94ER61937; the U.S. Environmental Protection Agency under grant XA-83600001-1; the National Science Foundation (NSF) through the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507 and the Penn State Center for Climate Risk Management; and other government, industry, and foundation sponsors of the Joint Program on the Science and Policy of Global Change. For a complete list of sponsors, please visit <http://globalchange.mit.edu/sponsors/all>.

REFERENCES

- Barforoushi, T., M.P. Moghaddam, M.H. Javidi, and M.K. Sheikh-El-Eslami (2010). "Evaluation of Regulatory Impacts on Dynamic Behavior of Investments in Electricity Markets: A New Hybrid DP/GAME Framework." *IEEE Transactions on Power Systems* 25 (4): 1978–1986. <https://doi.org/10.1109/TPWRS.2010.2049034>.
- Bellman, R.E. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University Press.
- Bergerson, J.A. and L.B. Lave (2007). "Baseload Coal Investment Decisions under Uncertain Carbon Legislation." *Environmental Science and Technology* 41(10): 3431–3436. <https://doi.org/10.1021/es062198e>.
- Bertsekas D., J.N. Tsitsiklis, (1996). *Neuro Dynamic Programming*. Belmont, MA: Athena Scientific.
- Birge, J.R. and F. Louveaux (2011). *Introduction to Stochastic Programming*. Verlag: Springer. <https://doi.org/10.1007/978-1-4614-0237-4>.
- Bistline, J.E. and J.P. Weyant (2013). "Electric sector investments under technological and policy-related uncertainties: a stochastic programming approach." *Climatic Change* 121:143–160. <https://doi.org/10.1007/s10584-013-0859-4>.
- Bosetti, V., C. Carraro, and E. Massetti (2009). "Banking Permits: Economic Efficiency and Distributional Effects." *Journal of Policy Modeling* 31(3): 381–403. <https://doi.org/10.1016/j.jpplmod.2008.12.005>.

- Chen, Y.-H., S. Paltsev, J.M. Reilly, J.F. Morris and M.H. Babiker (2016). "Long-term Economic Modeling for Climate Change Assessment." *Economic Modelling* 52(Part B): 867–883.
- Chronopoulos, M., V. Hagspiel and S.-E. Fleten (2016). "Stepwise Green Investment under Policy Uncertainty." *The Energy Journal* 37(4): 87–108. <https://doi.org/10.5547/01956574.37.4.mchr>.
- Crost, B. and C. Traeger (2010). *Risk and Aversion in the Integrated Assessment of Climate Change. CUDARE Working Paper No. 1104*.
- De Ghellinck G., (1960). Les problemes de decisions sequentielles. *Cahiers du Centre d'Etudes de Recherche Operationnelle* Vol. 2, pp. 161–179.
- De Jonghe, C., B.F. Hobbs, and R. Belmans (2012). "Optimal Generation Mix with Short-Term Demand Response and Wind Penetration." *IEEE Transactions on Power Systems* 27 (2): 830–839. <https://doi.org/10.1109/TPWRS.2011.2174257>.
- 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.
- Ehrenmann, A. and Y. Smeers (2011). "Generation Capacity Expansion in a Risky Environment: A Stochastic Equilibrium Analysis." *Operations Research* 59 (6): 1332–1346. <https://doi.org/10.1287/opre.1110.0992>.
- Energy Information Administration (EIA) (2015): *Annual Energy Outlook 2015: With Projections to 2040. U.S. Energy Information Administration, Washington, D.C.*
- Farias D.P., B. Van Roy, (2003). "The linear programming approach to Approximate Dynamic Programming." *Operations Research* 51(6): 850–865. <https://doi.org/10.1287/opre.51.6.850.24925>.
- 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. <https://doi.org/10.1016/j.enpol.2010.04.008>.
- 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. <https://doi.org/10.1146/annurev.energy.27.122001.083408>.
- Hammit, J.K., R.A. Lempert and M.E. Schlesinger (1992). "A Sequential-Decision Strategy for Abating Climate Change." *Nature* 357(6376): 315–318. <https://doi.org/10.1038/357315a0>.
- Hertel, T. (1997). *Global Trade Analysis: Modeling and Applications*. Cambridge, UK: Cambridge University Press.
- Hobbs, B.F. (1995). "Optimization methods for electric utility resource planning." *European Journal of Operational Research* 83: 1–20. [https://doi.org/10.1016/0377-2217\(94\)00190-N](https://doi.org/10.1016/0377-2217(94)00190-N).
- 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. <https://doi.org/10.1016/j.eneco.2006.05.014>.
- Jin, S., A. Botterud, and S.M. Ryan (2014). "Temporal Versus Stochastic Granularity in Thermal Generation Capacity Planning with Wind Power." *IEEE Transactions on Power Systems* 29 (5): 2033–2041. <https://doi.org/10.1109/TPWRS.2014.2299760>.
- Kelly, D. and C. Kolstad (1999). "Bayesian Learning, Growth, and Pollution." *Journal of Economic Dynamics and Control* 23(4): 491–518. [https://doi.org/10.1016/S0165-1889\(98\)00034-7](https://doi.org/10.1016/S0165-1889(98)00034-7).
- Kettunen, J., D. Bunn and W. Blyth (2011). "Investment Propensities under Carbon Policy Uncertainty." *The Energy Journal* 32(1): 77–118. <https://doi.org/10.5547/ISSN0195-6574-EJ-Vol32-No1-4>.
- 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. <https://doi.org/10.1006/jjeem.1996.0028>.
- 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. <https://doi.org/10.1504/ijgei.2000.004441>.
- 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*.
- Ma, J., V. Silva, R. Belhomme, D.S. Kirschen, and L.F. Ochoa (2013). "Evaluating and Planning Flexibility in Sustainable Power Systems." *IEEE Transactions on Sustainable Energy* 4 (1): 200–209. <https://doi.org/10.1109/TSSTE.2012.2212471>.
- Manne A.S., (1960). Linear programming and sequential decisions. *Management Science* 6(3): 259–267. <https://doi.org/10.1287/mnsc.6.3.259>. <https://doi.org/10.5547/ISSN0195-6574-EJ-Vol16-No4-1>.
- Manne, A.S. and R.G. Richels (1995). "The Greenhouse Debate: Economic Efficiency, Burden Sharing and Hedging Strategies." *The Energy Journal* 16(4): 1–37.
- Manne, A.S. and R.G. Richels (1992). *Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emission Limits*. Cambridge, Massachusetts: MIT Press.

- 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. <https://doi.org/10.1016/j.eneco.2004.04.026>.
- Mejia-Giraldo, D. and J.D. McCalley (2014a). "Maximizing Future Flexibility in Electric Generation Portfolios." *IEEE Transactions on Power Systems* 29 (1): 279–288. <https://doi.org/10.1109/TPWRS.2013.2280840>.
- Mejia-Giraldo, D. and J.D. McCalley (2014b). "Adjustable Decisions for Reducing the Price of Robustness of Capacity Expansion Planning." *IEEE Transactions on Power Systems* 29 (4): 1573–1582. <https://doi.org/10.1109/TPWRS.2013.2295166>.
- Morgan, G. and M. Henrion (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511840609>.
- Morris, J.F. (2013). *Electricity Generation and Emissions Reduction Decisions under Uncertainty: A General Equilibrium Analysis*. Massachusetts Institute of Technology Doctoral Dissertation, Engineering Systems Division.
- Morris, J., J. Reilly and H. Chen (2014). *Advanced Technologies in Energy-Economy Models for Climate Change Assessment*. MIT Joint Program on the Science and Policy of Global Change, Report 272, December.
- Munoz, F.D., B.F. Hobbs, J.L. Ho, and S. Kasina (2014). "An Engineering-Economic Approach to Transmission Planning Under Market and Regulatory Uncertainties: WECC Case Study." *IEEE Transactions on Power Systems* 29 (1): 307–317. <https://doi.org/10.1109/TPWRS.2013.2279654>.
- 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. <https://doi.org/10.5547/ISSN0195-6574-EJ-Vol18-No1-1>.
- Paltsev, S., J. M. Reilly, H.D. Jacoby, and J.F. Morris (2009). "The Cost of Climate Policy in the United States." *Energy Economics* 31(S2): S235-S243. <https://doi.org/10.1016/j.eneco.2009.06.005>.
- 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, August*.
- Park, H. and R. Baldick (2015). "Stochastic Generation Capacity Expansion Planning Reducing Greenhouse Gas Emissions." *IEEE Transactions on Power Systems* 30 (2): 1026–1034. <https://doi.org/10.1109/TPWRS.2014.2386872>.
- Powell, W.B. (2011). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. 2nd Ed. Hoboken, NJ: John Wiley & Sons, Inc. <https://doi.org/10.1002/9781118029176>.
- 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. <https://doi.org/10.1287/ijoc.1110.0470>.
- Rausch, S., G. Metcalf, J. Reilly, and S. Paltsev (2011). "Distributional Impacts of a U.S. Greenhouse Policy: A General Equilibrium Analysis of Carbon Pricing," in *U.S. Energy Tax Policy*, G. Metcalf (ed.), Cambridge: Cambridge University Press. pp. 52–107.
- 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). <https://doi.org/10.1021/es2034729>.
- Reinelt, P.S. and D.W. Keith (2007). "Carbon Capture Retrofits and the Cost of Regulatory Uncertainty." *The Energy Journal* 28(4): 101–128. <https://doi.org/10.5547/ISSN0195-6574-EJ-Vol28-No4-5>.
- 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. <https://doi.org/10.1016/j.eneco.2008.06.005>.
- Roh, J.H., M. Shahidehpour, and L. Wu (2009). "Market-Based Generation and Transmission Planning with Uncertainties." *IEEE Transactions on Power Systems* 24 (3): 1587–1598. <https://doi.org/10.1109/TPWRS.2009.2022982>.
- 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. <https://doi.org/10.1023/A:1008655831209>.
- 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*.
- Savage, S.L. (2009). *The Flaw of Averages: Why we Underestimate Risk in the Face of Uncertainty*. Hoboken, NJ: John Wiley & Sons, Inc.
- Schweitzer, P. and A. Seidman. (1985). Generalized polynomial approximations in Markovian decision processes. *J. of Math. Anal. and Appl.* 110: 568–582. [https://doi.org/10.1016/0022-247X\(85\)90317-8](https://doi.org/10.1016/0022-247X(85)90317-8).
- Shortt, A., J. Kilvoluma, and M. O'Malley (2013). "Accommodating Variability in Generation Planning." *IEEE Transactions on Power Systems* 28 (1): 158–169. <https://doi.org/10.1109/TPWRS.2012.2202925>.
- Stern, D. (2012). "Interfuel Substitution: A Meta-Analysis." *Journal of Economic Surveys* 26 (2): 307–331. <https://doi.org/10.1111/j.1467-6419.2010.00646.x>.
- USGCRP [U.S. Global Change Research Program] (2009). *Global Climate Change Impacts in the United States*. New York: Cambridge University Press.

- 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 (2012). "Analysis of climate policy targets under uncertainty." *Climatic Change* 112(3–4): 569–583. <https://doi.org/10.1007/s10584-011-0260-0>.
- 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, November.*
- 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. <https://doi.org/10.1007/s10584-008-9406-0>.
- Webster, M.D. (2002). "The Curious Role of Learning: Should We Wait for More Data?" *The Energy Journal* 23(2): 97–119. <https://doi.org/10.5547/ISSN0195-6574-EJ-Vol23-No2-4>.
- Wigley, T.M.L., R. Richels and J.A. Edmonds (1996). "Economic and Environmental Choices in the Stabilization of Atmospheric CO₂ Concentrations." *Nature* 379(6562): 240–243. <https://doi.org/10.1038/379240a0>.
- Wogrin, S., E. Centeno, and J. Barquin (2011). "Generation Capacity Expansion in Liberalized Electricity Markets: A Stochastic MPEC Approach." *IEEE Transactions on Power Systems* 26 (4): 2526–2532. <https://doi.org/10.1109/TPWRS.2011.2138728>.
- Yohe, G., N.G. Andronova and M.E. Schlesinger (2004). "To Hedge or Not Against an Uncertain Climate Future?" *Science* 306(5695): 416–417. <https://doi.org/10.1126/science.1101170>.