SMART: A Stochastic Multiscale Model for the Analysis of Energy Resources, Technology and Policy

Warren B. Powell*

Abraham George*
Jeff Stewart†

Alan Lamont†

*Department of Operations Research and Financial Engineering
Princeton University, Princeton, NJ 08544

†Science and Technology Directorate
Lawrence Livermore National Laboratory
Livermore, CA 94550

August 17, 2009

Abstract

We address the problem of modeling long-term energy policy and investment decisions while retaining the important ability to capture fine-grained variations in intermittent energy and demand, as well as storage. In addition, we wish to capture sources of uncertainty such as future energy policies, climate, and technological advances, in addition to the variability as well as uncertainty in wind energy, demands, prices and rainfall. Accurately modeling the value of all investments such as wind and solar requires handling fine-grained temporal variability and uncertainty in wind and solar, as well as the use of storage. We propose a modeling and algorithmic strategy based on the framework of approximate dynamic programming (ADP) that can model these problems at hourly time increments over a multidecade horizon, while still capturing different types of uncertainty. This paper describes initial proof of concept experiments for an ADP-based model, called SMART, by describing the modeling and algorithmic strategy, and providing comparisons against a deterministic benchmark as well as initial experiments on stochastic datasets.

1 Introduction

Over the next few decades, the United States needs to transition to an energy system with much lower carbon emissions and reduced dependence on unreliable supplies. Planning this transition has to be made in the presence of uncertainty surrounding technology, climate, government policy and commodity prices. This process requires significant investment in R&D and new infrastructure. The investments have to be made in the presence of multiple sources of uncertainty that could influence the future energy system, such as the mix of primary energy (nuclear, renewable, sequestered fossil), and the balance of energy carriers such as electricity, hydrogen, alcohol and synthetic fuels, and end-use strategies (level of demand, demand response). For example, should we begin to put into place carbon capture-ready coal power plants now, or should we continue with pulverized coal for the next few decades? Should we start to reduce emissions now using the technologies that we have? Or should we allow the carbon emissions to continue at the current rate for now and wait to reduce them to a greater degree later, when new technologies are developed?

The dynamics of our system are driven by exogenous factors (wind, solar, climate, technology, prices, supplies and demands) and market decisions (how much energy should be produced by each source, how much of this energy should be used for each type of demand, how much energy should be stored, and most important, how much capacity should be added or retired each year in response to changing conditions). To model investment decisions and government policies, the model has to extend over a multi-decade horizon. To accurately represent the economics of energy technologies, our model has to capture the hourly fluctuations of wind, daily solar cycles, seasonal climate changes, and annual changes in energy technologies. Finally, it is widely understood that a key enabling technology for intermittent energy is storage. In order to be able to calculate the amount of energy in each type of storage, we require a model that can step forward through time.

Past models have represented market investment decisions by assuming they can be reasonably approximated as optimal solutions to aggregate, deterministic linear programs. These models have been effective at capturing fundamental differences in supply and demand, such as the groundbreaking policy conclusions from the PIES model (Hogan (1975)). In recent years, there has been growing interest in the use of models to guide investment and policy decisions in new technologies such as wind and solar, and to guide government policies such as carbon taxes and research and development to meet goals such as 20 percent renewable by 2030. It has long been recognized, however, that models

that represent intermittent sources such as wind and solar overstate their value if they use averages (wind turbines are dramatically more productive if wind blows at a constant speed). The problem with intermittent sources is that they are not dispatchable, requiring investments in other energy technologies to make up for low periods, even if the average production is very high (see Lamont (2008) for a thorough discussion of these issues). Energy policy researchers have found that using time increments longer than an hour begin to introduce significant errors. Storage can overcome the limitations of intermittent sources, but representing storage imposes additional demands on a model.

Separate from the need to handle fine-grained time scales is the need to capture different types of uncertainty. This ranges from uncertainty about wind (wind operators often have to commit to a certain level of production a day in advance), to uncertainties about policies (imposition of a carbon tax), technology (will we be able to sequester carbon), prices (what will be the price of oil in five years) and climate (how will global warming affect snow accumulation and rainfall). These uncertainties can impact the risks incurred in an investment, significantly reducing the appeal of technologies that look attractive if the future unfolds as expected.

The goals in the development of our energy model include the following features: a) Stochastic - the model needs to handle a range of different forms of uncertainty, from hourly wind, solar and demands, to daily and seasonal variations in prices and rainfall, to longer term changes in technology, commodity supplies and prices, government policy and climate; b) Multiscale - the model has to be able to handle variations on multiple temporal scales (hourly, daily, seasonal and yearly time scales) as well as multiple spatial scales; c) Storage - the ability to store energy can have a significant impact on the benefits of intermittent sources of energy, as well as on the overall mix of energy technologies; d) Benchmarks - it is not hard to get a model that handles all these features and produces numbers, but it was a major goal of our research to be able to benchmark our numbers in some way to validate the accuracy of the algorithm.

This paper describes an early implementation of “SMART” (Stochastic Multiscale model for the Analysis of energy Resources, Technology and policy). SMART uses a control-theoretic framework to step forward through time in hourly increments over a multidecade horizon, which also makes it possible to model energy storage. Approximate dynamic programming is used to produce policies that properly model the flow of information when making decisions. Although designed to handle stochastic problems, SMART can be applied to deterministic problems, making it possible to compare the performance of the model to a deterministic linear programming model. We show that

SMART closely matches the optimal solution from a deterministic linear program, which provides an important benchmark for our ADP algorithm. We also show that it exhibits robust behaviors when applied to a stochastic problem. To the best of our knowledge, this is the first stochastic, multiscale model that has been proposed for energy policy planning which properly handles the flow of information.

SMART can be used for both long-term policy studies, as well as economic analyses of portfolios of energy production and storage technologies. It can be used to simulate energy investment decisions, under the assumption that a cost-minimizing model mimics market behavior, or as a normative model to determine how to reach specific energy targets. It can be used to quantify risk, and to study the impact of different energy investment strategies. It can also provide accurate estimates of the value of intermittent energy sources and the value of storage using realistic assumptions about our ability to forecast energy supply.

This paper is organized as follows. Section 2 reviews the literature on energy policy models. Section 3 provides a sketch of our modeling and algorithmic strategy. The full mathematical model is described in section 4, capturing both the physical dynamics and the flow of information and decisions, but without providing a method for making decisions. Section 5 provides two optimization formulations to drive the process of making decisions: a deterministic linear programming model (which we later use as a benchmark), and a stochastic optimization model. Section 6 describes an algorithmic strategy for solving the stochastic optimization problem based on approximate dynamic programming. In section 7, we describe benchmarking experiments which compare the solution obtained using approximate dynamic programming to the optimal solution produced by a linear programming algorithm using a deterministic model. Experiments on stochastic applications are also reported to demonstrate specific behaviors that we would expect from a nonanticipative decision rule. Section 8 concludes the paper.

2 Review of energy models

There are three bodies of literature that are relevant to our research: deterministic energy investment models, stochastic optimization models in energy, and storage models.

Following the original PIES model (Hogan (1975)), a series of deterministic optimization models have been developed to model energy investment decisions over multidecade horizons. These include

the National Energy Modeling System (NEMS) (see NEMS (2003)) that takes as input demands and costs of various energy resources from a family of models. The MARKet ALlocation (MARKAL) model is a deterministic, optimization-based system that integrates the supply and end-use sectors of an economy (Fishbone & Abilock (1981) and Loulou et al. (2004)). MARKAL minimizes the discounted total cost of the system which includes investment, maintenance and operations costs of the technologies and the purchase costs of fuels, and also accounts for revenue generated. META*Net (Lamont (1997), Lamont (2008)), developed at the Lawrence Livermore National Laboratory, combines the techniques used in NEMS as well as MARKAL in that it performs market equilibration in a linear programming-based model.

The early stochastic energy investment models can be divided into two categories. We refer to the first category as partially stochastic models, such as the stochastic version of MARKAL (Johnson et al. (2006)) developed by the Environmental Protection Agency. In this version of MARKAL, a “stochastic wrapper” is used, wherein random parameters are sampled from probability distributions, after which a deterministic optimization problem is solved. This process is repeated a number of times, after which results are averaged. This means that for each run on a sample realization, decisions are allowed to see events in the future. Averaging can produce bizarre behaviors such as building 60 percent of a nuclear power plant if we think there is a 60 percent probability of a carbon tax being implemented. Kanudia & Loulou (1998) presents a version of MARKAL which explicitly incorporates multiple scenarios and includes as many replications of the MARKAL variables as the number of scenarios considered. The optimization is done by taking into account the probability of occurrence of each scenario, the objective function being a sum of the individual scenario objective functions weighted by the corresponding probabilities. However, this version does not scale up to handle more than a few possible future scenarios.

Another example of a partially stochastic model is the stochastic MESSAGE (Model for Energy Supply Strategy Alternatives and their General Environmental Impact) from the International Institute for Applied Systems Analysis (IIASA). The conventional MESSAGE model (Messner & Strubegger (1995)) uses point estimates to forecast technology parameters. The stochastic MESSAGE model (Messner et al. (1996)) aims to reduce the risk of underestimating the total investment cost as a stochastic optimization problem, using a sufficiently large number of sample realizations of the investment costs.

The second category of stochastic models uses the framework of stochastic programming to handle

the modeling of random information and decisions (see Wallace & Fleten (2003) for an excellent survey of these models). Stochastic programming models use the notion of scenario trees to represent uncertain information. The challenge with this framework is that the number of scenario trees grows exponentially with the number of time periods, requiring the use of aggregation strategies such as large time steps and small numbers of scenarios per time period. See Birge & Louveaux (1997) for an-depth treatment of this approach, and Høyland & Wallace (2001), Dupacova et al. (2002) and Kaut & Wallace (2003) for research on generating scenario trees. A competing strategy uses dual information to generate cuts to approximate the impact of decisions now on the future. This was originally introduced as the L-shaped method Van Slyke & Wets (1969), and later adapted (for two-stage problems) using a Monte Carlo-based framework as stochastic decomposition (Higle & Sen (1991), Higle & Sen (1996)).

The last class of stochastic energy investment models can be viewed as pure simulation models. Perhaps the most visible example of this class is the Stochastic Energy Deployment Systems (SEDS) (see <http://seds.nrel.gov/>), which uses classical Monte Carlo simulation to step forward through time. Let ω^n represent the sample realization of all random variables for the n^{th} iteration of the model, and let $c_{te}(\omega^n)$ be the cost of energy technology $e \in \mathcal{E}$ in year t , while following sample path ω^n . SEDS makes additional investments, given by x_{te}^n , according to the expression

$$x_{te}^n = \beta_t^D \frac{\exp - \beta^s c_{te}(\omega^n)}{\sum_{e' \in \mathcal{E}} \exp - \beta^s c_{te'}(\omega^n)} \quad (1)$$

where β_t^D is a set of parameters (over time) which allow investments to be scaled to demand, and β^s is a scaling coefficient that controls the degree to which investments are directed toward lower cost technologies. There is no attempt to formally optimize energy investments. A major goal of the model is to provide fast estimates of the effects of policy changes in a fairly transparent way. The model works in yearly (or even multiyear) increments, so there is no attempt to capture fine-grained variations in wind and solar energy or demand.

There is a separate literature on the use of storage, largely divided between the management of water reservoirs and other forms of energy storage. The literature on water reservoir management is fairly extensive (see Fofoula-Georgiou & Kitanidis (1988), Lamond & Sobel (1995), Archibald et al. (1997) and Cervellera et al. (2005) for a sample). Algorithmic research addressing the problem of optimizing flow controls under uncertainty has progressed along two lines. The first uses dynamic programming, typically resorting to some sort of numerical or approximate strategy to handle the

curse of dimensionality that arises with multiple reservoirs (see Nandalal & Bogardi (2007) and the references cited there). A major breakthrough in this approach involves approximating the impact of decisions now on the future using cuts generated from the dual of the linear program in the next time period. Pereira & Pinto (1991) is the first to use this strategy in an energy setting under the name of stochastic, dual dynamic programming (SDDM), modeling a large, multireservoir hydroelectric system. A separate modeling and algorithmic strategy within the stochastic programming community uses the concept of scenario trees, where sample paths of exogenous information are generated, retaining the entire history at each node in the tree. Jacobs et al. (1995) and De Ladurantaye et al. (2007) give nice illustrations of this strategy in a hydroelectric setting. Heuristic approaches to reservoir management are described in Johnson et al. (1991), Edmunds & Bard (1990*b*) and Edmunds & Bard (1990*a*).

A separate literature has evolved around the modeling of storage devices to handle either structural or random variations between our ability to generate power and the demand for power. The recent paper Brunetto & Tina (2007) uses a deterministic model to solve the problem of day-ahead electricity markets, completely ignoring the high level of uncertainty that arises in this setting. Castronuovo & Lopes (2004) describes wind as a stochastic process, but then represents only the variability of wind, and not its uncertainty. Korpaas et al. (2004) and Paatero & Lund (2005) also use deterministic models to study storage for wind. Garcia-Gonzalez et al. (2008) and Brown et al. (2008) use the idea of stochastic wind scenarios, but allows decisions to “see” the wind in the future within a particular scenario (violating what are known as nonanticipativity conditions).

3 Strategy for a stochastic, multiscale model

We adopt a classical control-theoretic framework where dispatch, storage and investment decisions at hour h in year t are represented by the vector x_{th} , which is determined by a decision function (or *policy*, in the language of dynamic programming) $X^\pi(S_{th})$ which depends on the system state variable S_{th} . S_{th} is designed to include only the information available at time (t, h) , and as a result decisions are not allowed to anticipate events in the future.

Our strategy involves stepping forward through time as a simulation model would, but repeating the process while learning from each iteration. Once each year (corresponding to $h = 0$), we make investment decisions for new energy capacity that will arrive in a future year. Then, for the remainder

of the year we step forward solving hourly dispatch problems which determines how much energy to produce from each source to satisfy each type of demand. Energy demands are comprised of residential, commercial and industrial demands for electricity and natural gas (residential uses include cooking and heating), as well as freight transportation and aircraft demands for electricity, hydrogen and hydrocarbon fuels. Primary energy resources include coal, petroleum, natural gas, wind, solar, temperature gradients (geothermal, ocean thermal), nuclear, biomass and tides. These undergo conversion processes in plants to generate electricity, hydrocarbon fuels or hydrogen, each of which is then channeled to the appropriate sources of demand. While most of the conversion plants that handle carbon-based energy sources release CO_2 to the atmosphere, some technologies are able to sequester CO_2 . The dispatch problem also includes decisions of how much to store, and how much to convert to different types of energy (for example, natural gas can be used to generate electricity which in turn can be used to convert biomass to fuel or to compress hydrogen).

Figure 1 illustrates a portion of the hourly dispatch problem for the electricity market as it is represented in SMART. If there is a shortage of energy supplied from these resources, there also exists the option of importing electricity. Each node in the network represents an aggregated version of a facility in a typical energy network. Resource nodes provide the supply of raw resource (such as coal, natural gas, wind, solar, biomass, nuclear) at a certain cost. Conversion nodes take as inputs the raw resource and some type of power as input to generate more power in some form such as electricity, natural gas and transportation fuel. The flows out of each of these nodes are constrained by the capacity available at that node, which is controlled by investment decisions. Demand nodes capture the demand for one of the various types of power. Market nodes accumulate energy from different sources to be distributed to the various types of demand (these can be sectors or locations). We may also add other nodes in this illustration to capture costs involved in the transportation of the energy resource (for example, from the source to the conversion plant) and the distribution of the final output (for example, electricity through transmission lines).

Our modeling and algorithmic strategy involves simulating our way through the time periods, solving small optimization problems to handle dispatch and storage decisions (each hour) and investment decisions (each year). As we step forward through time, we use Monte Carlo simulation to sample realizations from any random variables. Our challenge is designing a decision function $X^\pi(S_{th})$ that produces economically rational decisions, without violating restrictions on the availability of information.

Figure 1: Reduced version of the hourly dispatch problem with electricity demands only.

We begin by providing an outline of the model of the problem in section 4 (the complete model is given in Appendix A.1), followed by both deterministic and stochastic optimization formulations in section 5.

4 Outline of the energy model

We use a classical modeling framework from control theory to represent our system. We let $t \in \mathcal{T}$ index years, and $0 \leq h < 8760$ index the hours within a year. Our model consists of five fundamental components:

- S_{th} = The state of the system at time (t, h) .
- x_{th} = The vector of decisions at time (t, h) .
- W_{th} = The vector of random variables representing all the new information that just became available in the hour prior to (t, h) .
- $S^M(S_{th}, x_{th}, W_{t,h+1})$ = the transition function (or system model) that governs the evolution from state S_{th} to $S_{t,h+1}$ (or from $S_{t,H}$ to $S_{t+1,0}$) given decision x_{th} and new information $W_{t,h+1}$.
- C_{th} = The contribution function capturing rewards earned during the hour preceding (t, h) .

The state variable consists of the current status of our energy investments R_{th} , energy held in

different forms of storage y_{th} , energy demands D_{th} , rainfall levels p_{th} , and a generic “state of the world” vector ρ_{th} that captures information about technology, climate, government policy, and the availability of intermittent energy such as wind and solar. Of particular interest are the vectors R_{th} and y_{th} since these constitute the state variables that we directly control through energy investments and our control of the flow of energy into storage. We let the attributes of energy resources or storage be represented by the vector a which can capture the type of technology, its location and age. We let $\mathcal{A}$ be the set of all possible attribute vectors, and we let $\mathcal{A}^{Conv}$ and $\mathcal{A}^{Stor}$ represent the sets of attribute vectors that corresponding to conversion and storage resources, respectively. Our primary energy investment decision is investment in conversion resources (energy generating facilities). We let R_{tha} be the level of investment in a technology with attribute $a \in \mathcal{A}$ (for example, this could be the megawatthours of generating capacity of a wind farm), and let $R_{th} = (R_{tha})_{a \in \mathcal{A}}$ be the vector of energy investments. If $a \in \mathcal{A}^{Stor}$ refers to storage such as a reservoir, an underground salt mine that holds compressed air, a battery or a flywheel, then y_{tha} is the amount of energy stored in this device at time (t, h) .

The decision variable x_{th} consists of the vector of investments in energy conversion capacity x_{th}^{cap} and dispatch decisions x_{th}^{disp} . We assume we make decisions on energy investments only once a year at hour $h = 0$, which means we are only interested in x_{t0}^{cap} , with $x_{th}^{cap} = 0$ for $h > 0$. Energy dispatch decisions govern the flow of energy from each source node (e.g. a coal mine or source of natural gas), through conversion nodes (coal, nuclear and oil plants), through various transmission and storage nodes, and finally to demand nodes. A portion of the dispatch network is shown in figure 1. This is the network at a point in time (a particular hour). These networks are linked over time solely through the deposits to and withdrawals from storage devices. We let a be the attributes of a node in the dispatch network, which would represent moving energy over a link in the network. $x_{th,aa'}$ is the amount of energy moved from node a to node a' . We let $\theta_{aa'}$ capture transmission and storage losses when moving from a to a' ; θ can also handle changes in units (e.g. from barrels of oil to megawatthours of energy).

When we make decisions, they are governed by physical constraints. For example, there are conservation of flow constraints in the dispatch network, as well as upper bounds on the amount of energy that can be stored in a particular device or location at a point in time. The flow conservation

constraints are given by

$$\sum_{a' \in \bar{\mathcal{A}}(a)} \theta_{a'a} x_{th,a'a}^{disp} - \sum_{a'' \in \bar{\mathcal{A}}(a)} x_{th,aa''}^{disp} = 0, \text{ for nonstorage locations,} \quad (2)$$

$$y_{tha} + \sum_{a' \in \bar{\mathcal{A}}(a)} \theta_{a'a} x_{th,a'a}^{disp} - \sum_{a'' \in \bar{\mathcal{A}}(a)} x_{th,aa''}^{disp} \geq 0. \text{ for storage locations.} \quad (3)$$

All decisions are limited by nonnegativity constraints. We let $\mathcal{X}_{th}$ be the feasible region at time (t, h) , and we require that $x_{th} \in \mathcal{X}_{th}$.

Once we have made a decision, the system then evolves over time, with new information arriving that also changes the state of the system. We view new information as a random variable which represents some sort of exogenous change to one of our state variables. Thus, $\hat{R}_{th}$ can be exogenous changes to our energy investments (Hurricane Katrina taking drilling rigs out of commission), $\hat{D}_{th}$ can be random changes in demand for different types of electricity, $\hat{p}_{th}$ can be the random change in the amount of rainfall, and $\hat{\rho}_{th}$ can be a change in technology, climate or policy (the government just imposed a carbon tax). These changes may take place on hourly time scales, but will often occur on longer time scales (we model technology as changing only once per year). We let W_{th} capture the vector of all these sources of random information that occurred in the hour preceding (t, h) .

Later, we need some language for describing a set of sample realizations. Let Ω be the set of all possible realizations of the sequence $(W_{11}, \dots, W_{1H}, W_{21}, \dots, W_{2H}, \dots, W_{TH})$ where T is the number of years and H is the number of time periods in a year. We let $W_{th}(\omega)$ be a sample realization of the random variables in W_{th} when we are following sample path ω . If we are following sample path ω^n for iteration n , we will index variables such as the state variable using S_{th}^n to represent the actual state at time (t, h) while following sample path ω^n .

All of the physics of the problem are captured by the transition function (also known as the state model, plant model or simply model) using

$$\begin{aligned} S_{t,h+1} &= S^M(S_{th}, x_{th}, W_{t,h+1}), \quad h = 0, 1, \dots, H-1, \\ S_{t+1,0} &= S^M(S_{tH}, x_{tH}, W_{t+1,0}). \end{aligned}$$

For our model, the transition function primarily captures the random changes. For example, changes

in R_{th} and ρ_{th} would be written

$$R_{t,h+1} = R_{th} + x_{th}^{cap} + \hat{R}_{t,h+1}, \quad (4)$$

$$\rho_{t,h+1} = \rho_{th} + \hat{\rho}_{t,h+1}. \quad (5)$$

We measure our performance by adding up costs over time. Total capacity acquisition costs, dispatch costs, fuel costs and operating costs are given by

$$C_{th}^{cap}(S_{th}, x_{th}^{cap}) = \text{The incremental capital costs in hour } h \text{ of year } t,$$

$$C_{th}^{disp}(S_{th}, x_{th}^{disp}) = \text{The total costs resulting from the dispatch of energy in hour } h \text{ of year } t.$$

We may express the total cost function in hour h of year t using,

$$C_{th}(S_{th}, x_{th}) = C_{th}^{cap}(S_{th}, x_{th}^{cap}) + C_{th}^{disp}(S_{th}, x_{th}^{disp}).$$

We have not addressed the problem of how to make decisions. In the next section, we present two models, a deterministic linear programming model, and a stochastic model, which lead to different methods for making a decision.

5 Optimization formulations

Our modeling framework above left unanswered the mechanism for actually making decisions. Section 5.1 describes a deterministic linear programming formulation that can be solved using commercial packages. Section 5.2 provides a stochastic optimization formulation, for which exact solution algorithms are not available. The remainder of the paper then focuses on designing an approximation strategy to solve the stochastic optimization model.

5.1 A deterministic linear programming model

If we assume that all the exogenous information is deterministic (and known in advance), we can formulate the decision problem as a single (albeit very large) linear program. The objective function is given by

$$\min_x \sum_{t \in \mathcal{T}} \sum_{h \in \mathcal{H}} C_{th}(S_{th}, x_{th}).$$

This has to be solved subject to two sets of constraints. The first set governs decisions made at a point in time, which is given by equations (15) - (20) (in appendix A.1). The second set of equations are the transition equations that link activities over time (basically the transition function), which are captured by equations (21) - (25) (in appendix A.1). It is significant that the linear programming model views all these constraints together, over all points in time. This is why a linear programming model (even with deterministic formulations) quickly becomes too large to solve. For example, a five-year, spatially aggregate model with hourly dispatch decisions produced a constraint matrix with over 90 million rows and columns.

To hold down the size of the model, it is possible to discretize into coarser increments than hourly. However, in this case the input of intermittent sources (represented by ρ_{tha}^{exo}) has to be averaged over longer periods. Wind and solar appear more attractive when they are averaged over longer periods of time.

5.2 A stochastic optimization model

The linear programming model suffers from two important limitations. First, it is unable to handle longer planning horizons while retaining the ability to model intermittent supplies and demand on an hourly level (along with storage). Second, it cannot handle the different forms of uncertainty that arise in energy policy analysis. For this reason, we pose the problem of controlling the stochastic version of the problem in terms of finding a policy (or decision function) to make decisions about dispatch, storage and investments over time. We represent the decision function using

$$X_{th}^{\pi}(S_{th}) = \text{A function that returns a decision } x_{th} \in \mathcal{X}_{th} \text{ given the state } S_{th} \text{ under the policy } \pi \in \Pi.$$

The set Π refers to the set of potential decision functions, whose meaning becomes clearer below. At this point, we refer abstractly to the space of all possible decision functions.

The problem being stochastic, our goal is to find the optimal policy from among the set Π of policies by solving

$$\min_{\pi \in \Pi} \mathbb{E} \left\{ \sum_{t \in T} \sum_{h \in \mathcal{H}} C_{th}(S_{th}, X_{th}^{\pi}(S_{th})) \right\}, \quad (6)$$

where $S_{t,h+1} = S^M(S_{th}, x_{th}, W_{t,h+1})$ (for $h = 0, \dots, H-1$) and $S_{t+1,0} = S^M(S_{tH}, x_{tH}, W_{t+1,0})$. The challenge, of course, is designing the policy $X_{th}^{\pi}(S_{th})$.

We note that by construction of our decision function (in particular, its dependence on the state variable) it is not allowed to depend on events in the future. This seemingly essential property is violated by the stochastic models that use “LP-averaging,” where random variables are sampled over all time periods after which a single linear program is solved.

6 Algorithmic strategy

There are a number of strategies that have been used to solve stochastic optimization problems as stated in equation (6). These include various types of myopic policies (such as the logit investment rule in SEDS), rolling horizon policies (which use a point forecast of the future to make decisions now), simulation-optimization where we optimize the parameters of a myopic policy (Swisher et al. (2000)), stochastic programming (Higle & Sen (1996), Birge & Louveaux (1997), Wallace & Fleten (2003)), and dynamic programming. We adopt the framework of dynamic programming in section 6.1, which introduces a number of computational hurdles. These are addressed in section 6.2 using approximate dynamic programming.

6.1 Dynamic Programming Formulation

Our strategy to solve the energy model involves formulating the model as a dynamic program, solving decision problems for each hour $h \in \mathcal{H}$ in each year $t \in \mathcal{T}$. We first define the value $V_{th}(S_{th})$ of a state S_{th} as the sum of the contributions that we expect to make, starting in state S_{th} , if we act optimally (make optimal decisions) until the end of the time horizon. Bellman’s equation enables us to recursively compute the optimal value functions associated with each state,

$$V_{th}(S_{th}) = \max_{x_{th} \in \mathcal{X}_{th}^{disp,n}} (-C_{th}(S_{th}, x_{th}) + \mathbb{E}(V_{t,h+1}(S_{t,h+1}) | S_{th})), \quad (7)$$

where $S_{t,h+1} = S^M(S_{th}, x_{th}, W_{t,h+1})$.

As pointed out earlier, for a given year t , we assume that capacity additions are made only at the beginning ($h = 0$) of the year with new capacity arriving in some future year; the capacity resource state remains unchanged during the remaining hours. The value function $V_{th}(S_{th})$ in equation (7) captures the information about the future state of the system while solving the problem for the current time, including the value of water left behind in the reservoir after the decisions in the current time period are implemented.

In the remainder of the paper we consider only a single form of storage, namely the reservoir that stores water collected from precipitation. This means that the set $\mathcal{A}^{Stor}$ consists of a single aggregate reservoir, but it is straightforward to adapt our model to multiple reservoirs, as well as multiple types of storage. The water from the reservoir may be used to satisfy the energy demands now or stored to be used to satisfy demands later in time. Hydro-energy is typically among the cheaper types of energy and tends to be used before more expensive energy types (such as coal or nuclear) if we do not provide for storage. Ideally, our model of reservoir management should mimic real-life behavior where water would be allowed to collect and be stored in the reservoir during periods of heavy precipitation to be used later during dry periods.

Solving the dynamic program in equation (7) is generally hard even for problems with small state spaces. The traditional dynamic programming approach for solving equation (7) suffers from three curses of dimensionality arising from the state space ($\{S_{th}\}_{h \in \mathcal{H}, t \in \mathcal{T}}$), the exogenous information ($\{W_{th}\}_{h \in \mathcal{H}, t \in \mathcal{T}}$), and the action space ($\{\mathcal{X}_{th}^{disp}\}_{h \in \mathcal{H}, t \in \mathcal{T}}$) (see Powell (2007), chapter 4). We now describe the techniques we use to overcome these.

We first define the post-decision state variable, $S_{th}^x = S^{M,x}(S_{th}, x_{th})$. The post-decision state for hour h is the state of the system immediately after making the decisions at hour h , but before the occurrence of the random events ($W_{t,h+1}$) between hours h and $h + 1$. In our analysis, we focus on two important components of the state variable, namely the capacity resource state R_{th} and the reservoir level y_{th} . The post-decision transitions of these components can be written as,

$$\begin{aligned} R_{tha}^x &= R_{tha} + x_{tha}^{cap}, \quad a \in \mathcal{A}^{conv}, \\ y_{tha}^x &= y_{tha} + \sum_{a' \in \bar{\mathcal{A}}(a)} \theta_{a'a} x_{th,a'a}^{disp}, \quad a \in \mathcal{A}^{Stor}. \end{aligned}$$

We then construct the value function around the post-decision state $S_{th}^x = (R_{th}^x, y_{th}^x)$ using

$$V_{th}^x(S_{th}^x) = \mathbb{E}[V_{t,h+1}(S_{t,h+1}) | S_{th}^x]. \quad (8)$$

This allows us to avoid having to compute an expectation within the optimization formulation of equation (7). The modified dynamic program can be written as

$$V_{th}(S_{th}) = \max_{x_{th} \in \mathcal{X}_{th}^{disp}} (-C(S_{th}, x_{th}) + V_{th}^x(S_{th}^x)).$$

This leaves the problem of computing $V_{th}^x(S_{th}^x)$. For this, we turn to the techniques of approximate dynamic programming. It is important to realize that our strategy represents a significant departure from classical modeling and algorithmic strategies used in the approximate dynamic programming literature (notably Bertsekas & Tsitsiklis (1996) and Sutton & Barto (1998)) which approximate the value function around the pre-decision state. This strategy still leaves us with the challenge of computing or approximating the expectation. We avoid this by directly approximating the value function around the post-decision state.

6.2 An Approximate Dynamic Programming Algorithm

Our solution strategy involves choosing an appropriate functional form for the value functions. Our idea is very similar to that used in Pereira & Pinto (1991) for water reservoir management, with the significant difference that we are using separable, piecewise linear approximations of the value of resources in the future, where resources can be the amount of energy in storage as well as energy investments. This contrasts with the use in Pereira & Pinto (1991) of multidimensional cuts based on Benders decomposition, which carries with it provable convergence properties (Higle & Sen (1991), Ruszczyński (1993)). We have compared both of these strategies and found that the separable approximations produced much faster convergence than approximations based on Benders cuts (see Powell et al. (2004) for an analysis in the context of two-stage problems) and very accurate solutions for complex multistage problems that arise in transportation (see Powell & Topaloglu (2003), Powell & Topaloglu (2004), Topaloglu & Powell (2006)). However, both strategies use the same basic idea of forming piecewise linear approximations based on dual information.

Starting with initial value function approximations, we step forward in time, solving the decision problem at each time using the value function approximation of the future states. We repeat this procedure over several iterations, updating the value function approximations after each iteration. This iterative updating enables us to form an approximation of the expectation in equation (8) (actually, we only need to approximate the slopes). While this algorithmic strategy is not generally optimal, there are convergence proofs for special versions of the problem (discussed below), and in section 7 we benchmark against the linear programming formulation for a deterministic version of the problem which demonstrates very good agreement.

Using some policy π , we may write the objective function for time t as follows,

$$\begin{aligned} F_{th}^\pi(S_{th}) &= -C_{th}(S_{th}, X_{th}^\pi) - \sum_{h' > h} C_{th'}(S_{th'}, X_{th'}^\pi) - \sum_{t' > t} \sum_{h \in \mathcal{H}} C_{t'h}(S_{t'h}, X_{t'h}^\pi) \\ &= -C_{th}(S_{th}, X_{th}^\pi) + V_{th}^x(S_{th}^x), \end{aligned} \quad (9)$$

where $V_{th}^x(S_{th}^x) = -\sum_{h' > h} C_{th'}(S_{th'}, X_{th'}^\pi) - \sum_{t' > t} \sum_{h \in \mathcal{H}} C_{t'h}(S_{t'h}, X_{t'h}^\pi)$ is the same as the post-decision value function in equation (8). Since this function is unknown, we use an approximation which we denote as $\bar{V}_{th}^x(S_{th}^x)$. We now introduce the approximation that the value function is separable in the post-decision capacity resource state R_{th}^x and reservoir level y_{th}^x . We express this approximation as the sum of components denoting the value of the capacity resource state and the value of water remaining in the reservoir,

$$\begin{aligned} \bar{V}_{th}^x(S_{th}^x) &= \bar{V}_{th}^{cap}(R_{th}^x) + \bar{V}_{th}^{hydro}(y_{th}^x) \\ &= \sum_{a \in \mathcal{A}^{Conv}} \bar{V}_{tha}^{cap}(R_{tha}^x) + \bar{V}_{th}^{hydro}(y_{th}^x). \end{aligned}$$

It is easy to show that $V(S)$ is concave in R and y , and we approximate $V(R)$ and $V(y)$ using separable, piecewise linear functions (using a very fine discretization, since the quantities are continuous). It is important to recognize that we are only concerned about the derivative of $\bar{V}_{th}^x(S_{th}^x)$ rather than the actual value.

The hourly dispatch problem involves solving the optimization problem

$$x_{th}^{disp,n} = \arg \max_{x_{th} \in \mathcal{X}_{th}^{disp,n}} (-C_{th}^{disp}(S_{th}^n, x_{th}) + \bar{V}_{th}^{hydro}(y_{th}^x)) \quad (10)$$

where S_{th}^n is the specific state that we are in at hour h , year t , while following sample path ω^n . y_{th}^x is a function of both S_{th}^n and x_{th}^{disp} . Equation (10) defines our policy π , and the space of policies Π is the set of all possible value function approximations. The dispatch problem at time t is illustrated in figure 2, where the value function approximation (the value of water stored in the reservoir) is given by a piecewise linear function. It is important to emphasize that this is a very small linear program (for a spatially aggregate model), which scales fairly easily to spatially disaggregate applications (see Powell et al. (2002), Powell & Topaloglu (2005), Topaloglu & Powell (2006), Simao et al. (2009) for numerous transportation applications). This is important, since we have to solve 8,760 of these for each year of planning.

Figure 2: The hourly optimization problem with value function approximations for the reservoir level.

The annual capacity acquisition problem, solved at the beginning of each year, is given by

$$x_{t0}^{cap,n} = \arg \max_{x_{t0}^{cap}} \left(-C_{t0}^{cap}(S_{t0}^n, x_{t0}) + \bar{V}_{t0}^{cap}(R_{t0}^x) \right). \quad (11)$$

subject to

$$R_{t0}^x = R_{t0} + x_{t0}^{cap,n}, \quad (12)$$

$$x_{t0}^{cap,n} \geq 0. \quad (13)$$

This problem is illustrated in figure 3. Here, we use a value function approximation for each type of energy resource. The only decision of how much to purchase is handled by the piecewise linear value function giving the value of resources in the future.

Using ADP, we learn the value function approximations using iterative updating operations. Let n denote the iteration counter while following sample path ω^n . For resource allocation problems, we have found that it is effective to compute the marginal value of an additional unit of a resource (this

Figure 3: The annual capacity acquisition problem with value function approximations for the future capacity resource state.

can be energy resources R_{th} , or the energy in storage y_{th}). Let

- $\hat{v}_{tha}^{cap,n}$ = A sample estimate of the marginal value of increasing R_{tha} by one unit computed while following sample path ω^n ,
- $\hat{v}_{tha}^{hydro,n}$ = A sample estimate of the marginal value of increasing y_{tha} by one unit computed while following sample path ω^n .

$\hat{v}_{tha}^{hydro,n}$ is computed relatively simply by finding the marginal impact of the value of additional energy in storage for a particular hour. Computing $\hat{v}_{tha}^{cap,n}$ requires calculating the incremental value of additional energy investment over an entire year. This requires finding the incremental value of, say, additional wind turbine generating capacity for each hour over the year, taking into account the presence of storage. This calculation allows us to capture the impact of hourly, daily and seasonal variability, as well as the marginal impact of additional storage.

In SMART, we used the updating equations given in Powell (2007), section 11.3. Note that the value $\hat{v}_{th}^{hydro,n}$ or $\hat{v}_{th}^{cap,n}$ found at time t , iteration n , is used to update the value function approximation around the post-decision state variable in the *previous* time period (see Powell (2007), section 4.4).

We represent general updating operations for the value functions as follows,

$$\begin{aligned}\bar{V}_{t,h-1}^{hydro,n}(y_{t,h-1}^x) &\leftarrow U^V \left(\bar{V}_{t,h-1}^{hydro,n-1}, y_{t,h-1}^{x,n}, \hat{v}_{th}^{hydro,n} \right), \\ \bar{V}_{t,h-1}^{cap,n}(R_{t,h-1}^x) &\leftarrow U^V \left(\bar{V}_{t,h-1}^{cap,n-1}, R_{t,h-1}^{x,n}, \hat{v}_{th}^{cap,n} \right).\end{aligned}$$

The details of the calculation of the derivatives $\hat{v}_{th}^{hydro,n}$ and $\hat{v}_{th}^{cap,n}$ are somewhat involved and are left to appendices A.2 and A.3, respectively.

In figure 4, we provide an outline of the ADP algorithm to solve the energy model.

6.3 Discussion

The essential feature of the algorithm is that it steps forward through time, using a Monte Carlo sample of any random variable, solving sequences of very small linear programs. In this sense, the model scales well to handle more complex information processes, or a spatial representation of the problem (for example, dividing the country into 10 to perhaps 100 regions). If the model is being used over a 30 year horizon, in hourly increments, we would have to model over 250,000 time periods. This is quite manageable for an aggregate model (as described in our experimental work), but CPU times would start to become an issue with a sufficiently fine-grained spatial model.

We have illustrated the logic using value function approximations which ignore other aspects of the state variable such as the state of climate (is global warming progressing faster than expected), technology (was there a significant change in the performance of batteries or the cost of extracting and sequestering carbon), markets (is the price of oil over \$150) and government policy (is there a carbon tax).

We can handle these dimensions, in theory, by simply replacing the separable piecewise linear value functions $\bar{V}_{th}^{cap}(R_{th}^x)$ and $\bar{V}_{th}^{hydro}(y_{th}^x)$ with functions that also depend on the state variable ρ_{th} , as in $\bar{V}_{th}^{cap}(R_{th}^x|\rho_{th})$ and $\bar{V}_{th}^{hydro}(y_{th}^x|\rho_{th})$. The algorithm would be virtually unchanged (see the storage problem described in Nascimento & Powell (2009b) for an illustration). The vector ρ_{th} evolves exogenously over time, and is already being computed. All that is needed is that we have to index the value functions by this information.

Of course, this is much easier to write than to implement. The number of possible values of ρ_{th} is extremely large, and as a result finding a value function approximation $\bar{V}_{th}^{cap}(R_{th}^x|\rho_{th})$ for each ρ_{th}

Step 0: Set the iteration counter, $n = 1$. Initialize the value function approximations $\{\bar{V}_{t0}^{cap,0}(R_{t0}^x)\}_{t \in \mathcal{T}}$ and $\{\bar{V}_{th}^{hydro,0}(y_{th}^x)\}_{h \in \mathcal{H}, t \in \mathcal{T}}$. Initialize the state of the system, S_{0a}^0 . Set $t = 0$.

Step 1: Solve the annual problem of capacity addition:

$$F_{t0}^{cap}(S_{t0}^n) = \max_{x_{t0}^{cap} \geq 0} (-C^{cap}(S_{t0}^n, x_{t0}^{cap}) + \bar{V}_{t0}^{cap,n-1}(R_{t0}^x)) \quad (14)$$

We let $x_{t0}^{cap,n}$ denote the argument that solves equation (14) and $\hat{v}_{t0}^{+,n} = (\hat{v}_{t0a}^{+,n})_{a \in \mathcal{A}}$, the vector of numerical derivatives where $\hat{v}_{t0a}^{+,n} = [F_{t0}^{cap}(S_{t0}^n, R_{t0a}^n + 1) - F_{t0}^{cap}(S_{t0}^n, R_{t0a}^n)] \quad \forall a \in \mathcal{A}$.

Step 2: Set the hour, $h = 0$.

Step 2a: Observe $W_t(\omega^n) = (\hat{R}_t, \hat{D}_t, \hat{\rho}_t, \hat{p}_t)$.

Step 2b: Solve the hourly optimization model (equation (10)) to obtain $F_{th}^{disp}(S_{th}^n)$. Let $\nu_{th}^{disp,n} = (\nu_{tha}^{disp,n})_{a \in \mathcal{A}}$ be the vector of dual values corresponding to the flow conservation constraints represented by equation (3) and let $\hat{v}_{th}^{hydro,n}$ be the numerical derivative $[F_{th}^{disp}(S_{th}^n, y_{th}^n + 1) - F_{th}^{disp}(S_{th}^n, y_{th}^n)]$.

Step 2c: Sample the incoming precipitation data, $\hat{p}_{t,h+1}$ in order to compute the new reservoir level using equation (25).

Step 2d: Set $h = h + 1$. If $h < H$, go to **step 2**. Otherwise, continue.

Step 2e: Update the value function approximations for hydro-storage,

$$\bar{V}_{t,h-1}^{hydro,n}(y_{t,h-1}^x) \leftarrow U^V(\bar{V}_{t,h-1}^{hydro,n-1}, y_{t,h-1}^x, \hat{v}_{th}^{hydro,n}) \quad \forall h.$$

Step 2f: Compute the hourly state transitions:

$$S_{t,h+1}^n = S^M(S_{th}^n, x_{th}^n, W_{t,h+1}^n)$$

Step 3: Compute the combined dual values (see (33))

$$\hat{v}_{ta}^{cap,n} = \hat{v}_{t0a}^{+,n} + \sum_{h \in \mathcal{H}} \frac{\nu_{tha}^{disp,n}}{\hat{\rho}_{tha}^{exo}} \quad \forall a \in \mathcal{A}.$$

Step 4: Update the value function approximations for the capacity acquisitions,

$$\bar{V}_{t-1,0}^{cap,n}(R_{t-1,0}^x) \leftarrow U^V(\bar{V}_{t-1,0}^{cap,n-1}, R_{t-1,0}^x, \hat{v}_t^{cap,n}).$$

Step 5: Compute the state transitions:

$$\begin{aligned} S_{t,h+1}^n &= S^M(S_{th}^n, x_{th}^n, W_{t,h+1}^n) \\ S_{t+1,0}^n &= S^M(S_{tH}^n, x_{tH}^n, W_{t+1,0}^n). \end{aligned}$$

Step 6: Set $t = t + 1$. If $t < T$, go to **step 1**. Otherwise, continue.

Step 7: Set $n = n + 1$. If $n < N$, go to **step 1**. Otherwise, return the value functions $\{\bar{V}_{t0}^{cap,n}(R_{t0}^x)\}_{t \in \mathcal{T}}$ and $\{\bar{V}_{th}^{hydro,n}(y_{th}^x)\}_{h \in \mathcal{H}, t \in \mathcal{T}}$, which define the policy for adding capacity over the years $t \in \mathcal{T}$.

Figure 4: The ADP algorithm to solve the energy resource management problem.

is computationally impossible when ρ_{th} has more than two or three dimensions. There is a wide range of statistical strategies for handling this problem (see, for example, Hastie et al. (2001)) which represent promising avenues for research. There is also a growing base of literature in approximate dynamic programming that draws on techniques such as kernel regression (Ormoneit & Sen (2002), Fan & Gijbels (1996)) which offer considerable promise.

6.4 Convergence analysis

Our solution strategy is not optimal, but we offer both theoretical arguments and, in section 7, empirical evidence to support the claim that the quality of the solution is extremely good.

First, the algorithm for optimizing storage over a single year is an instance of a provably convergent algorithm, as given in Nascimento & Powell (2009a) (building on the same proof technique as that given in Nascimento & Powell (2009b)). This proof technique applies even if we index the value functions on the exogenous state vector ρ_{th} . Central to this convergence proof is that the storage (which is the only means by which one subproblem interacts with another) is a scalar (the amount stored in a reservoir). Just the same, optimizing the storage in water reservoirs is a particularly difficult problem, since we need to decide whether to store additional water in the rainy months of February and March to satisfy needs later in the summer, thousands of hours into the future.

Second, the use of separable, piecewise linear approximations has been proven to be optimal for two-stage problems if the second stage is separable, and near-optimal even when the second stage is nonseparable (see Powell et al. (2004)). Separable approximations are not, in general, optimal for nonseparable, multistage resource allocation problems. However, we can obtain provably convergent algorithms for multistage problems. If we were to introduce other forms of storage, the algorithm would only be an approximation due to our use of separable functions. But for practical uses, separable approximations have been found to scale very well to high dimensional problems with fast convergence.

Theoretical convergence proofs, while reassuring, are no guarantee of empirical performance. Our research has been preceded by extensive experimental research with algorithms of this type for transportation applications (see Godfrey & Powell (2002), Topaloglu & Powell (2006)). In the next section, we report on experimental comparisons of the approximate dynamic programming algorithm on deterministic and stochastic versions of the problem.

7 Experiments

In this section, we design experiments that focus on the ability of the ADP algorithm in section 6.2 to learn near-optimal policies in terms of how much to store, and what investment decisions to make. We begin by comparing ADP to LP for a deterministic version of the problem, and later

Figure 5: Time to solve the energy model as a function of the number of years in the simulation model. With increasing number of years in the model, runtimes increase dramatically in order to obtain an optimal solution while ADP produces near-optimal results in a fraction of that time.

address the stochastic case. If we design the problem of capacity additions over multiple years and the hourly dispatch of power as a single optimization problem as presented in section 5.1, we can obtain an exact optimal solution in a reasonable amount of computational time provided the planning horizon is limited to a small number of years. However, the time required to solve the LP increases dramatically as the number of years in the simulation horizon is increased. Figure 5 shows the computational time in hours to solve the linear program as a function of the number of years. We also show the time taken by an approximate dynamic programming algorithm to solve the corresponding models to within a small fraction of a percentage error (0.02 percent) with respect to the optimal solution. The computational memory requirements are also seen to be comparatively much larger while solving the LP model to optimality. In the experiments to follow, we will be comparing the ADP solutions to the optimal solutions provided by the LP technique.

We begin, in section 7.1, by describing the performance indices that we use to evaluate our algorithmic strategy. In section 7.2, we describe our experiments in a purely deterministic setting where we run tests over a single year planning horizon with deterministic precipitation. In section 7.3, we discuss experiments with stochastic precipitation scenarios. In section 7.4, we focus on experiments on capacity acquisition in a multi-decade setting. Finally, in section 7.5, we analyze the capacity acquisition results in the presence of an uncertain future.

7.1 Performance measures

In this section, we design performance measures to determine how closely the ADP algorithm matches the optimal solutions generated by an LP solver. We use three types of policies to construct our performance measures: ADP (the solution produced using approximate dynamic programming), OPT (the optimal solution of the deterministic problem produced by solving the linear program), and MYOPIC, which is the results of a myopic policy (the same as ADP with value function approximations set to zero). The myopic policy is used only to compute one of the performance metrics. Below, we let strategy s refer to ADP, OPT or MYOPIC.

We define the following terms that will be used for computing the statistics that we measure:

- F^s = The objective function value as computed using strategy s . For example, F^{MYOPIC} denotes the myopic solution and F^{OPT} represents the optimal objective function value.
- x_{tha}^s = The total usage corresponding to resource node with attribute vector $a \in \mathcal{A}^{Res}$ ($a^{Type} \in \{\text{coal, solar, nuclear, } \dots\}$) in hour h of year t according to strategy s ,

$$= \sum_{a' \in \vec{\mathcal{A}}(a)} x_{th,aa'}^s.$$
- y_{th}^s = The reservoir level in hour h of year t obtained by following strategy s .
- c_a = Cost of adding a unit of capacity of type a .
- x_{ta} = Amount of capacity type a added in year t using ADP.
- x_{ta}^* = The optimal amount of capacity type a to be added in year t .
- z_a^* = The optimal amount of capacity type a that is used in the starting year.

We use these to compute the following measures:

General performance indices:

- E_1^s = The error in the objective function (using strategy s) from the optimal objective function as a percentage of the optimal.

$$= \frac{F^s - F^{OPT}}{F^{OPT}} \times 100\%$$
- E_2^s = The error in the objective function (using strategy s) as a percentage of the error in the objective function for a myopic policy.

$$= \frac{E_1^s}{E_1^{MYOPIC}} \times 100\%$$

Performance indices for hydro-storage:

$$R_1^s = \frac{\sum_{t,h,a} |x_{tha}^s - x_{tha}^{OPT}|}{\sum_{t,h,a} |x_{tha}^{OPT}|}$$

$$R_2^s = \frac{\sum_{t,h} (y_{th}^s - y_{th}^{OPT})^2}{\sum_{t,h} (y_{th}^{MYOPIC} - y_{th}^{OPT})^2}$$

Indices for capacity-matching:

$$R_1^{cap} = \frac{\sum_a c_a \sum_t |x_{ta} - x_{ta}^*|}{T \sum_a c_a z_a^*} \times 100\%$$

$$R_2^{cap} = \frac{\sum_a c_a |\sum_t x_{ta} - \sum_t x_{ta}^*|}{T \sum_a c_a z_a^*} \times 100\%$$

E_2^s is a good indicator of the improvement in the solution quality obtained by using strategy s compared to a myopic solution. R_1^s measures the deviation of the solution from the optimal as a fraction of the optimal solution. R_2^s is a measure of how well strategy s improves over a myopic solution of the storage problem.

7.2 Experiments on hydro-storage

The first set of experiments focuses on testing how well the hydro-usage behavior of the solution from the ADP algorithm matches the optimal solution from the LP model, under the assumption of deterministic precipitation. We use a time horizon of a single year in order to avoid the additional complexities introduced as a result of the capacity acquisition part of the energy model.

We conducted our experiments in a variety of deterministic settings. The demands were (1) constant, (2) sinusoidally varying, or, (3) varying with no apparent pattern. Another setting that could be varied was whether or not wind (varying by the hour) was included as a competing source of energy. We tested our algorithm for a one-year horizon with storage value functions computed for every 10 hour increment. We then repeated the test for one hour increments. For some problems, this means that the decision of how much water to draw from the reservoir in March had to reflect usage behavior in July, thousands of hours into the future.

Figure 6 shows the reservoir level over the course of a one year period for a deterministic problem produced by solving the linear program which gives an optimal solution (figures (a) and (c)), and from the ADP algorithm (figures (b) and (d)). These were run for constant demand (figures (a)

(a) Optimal-constant demand

(b) ADP-constant demand

(c) Optimal-sinusoidal demand

(d) ADP-sinusoidal demand

Figure 6: Optimal reservoir levels (a) and (c) and reservoir levels from ADP algorithm (b) and (d) for constant and sinusoidal demand.

and (b)) and a sinusoidal demand (figures (c) and (d)). Note that the heaviest precipitation occurs during late winter and spring. The model used a slight penalty for holding water in the reservoir so that the optimal solution was to build up reserves as late as possible (this choice was made purely for algorithmic testing purposes). With this setup, the algorithm had to realize that it was necessary to begin building up reservoirs in the March-April time frame to meet needs that would occur thousands of hours into the future. The results show an extremely close match between the optimal solution and that produced by the ADP algorithm.

In table 1, we summarize the performance statistics of the ADP algorithm. We observe that the ADP algorithm is able to produce solutions that are consistently within a very small margin of error of the optimal solutions.

7.3 Hydro-storage with stochastic precipitation

We now consider the case where the precipitation is stochastic. We generate 50 different precipitation scenarios for the year, and use these to train the value function approximations for approximate

Dataset	Demand type	H (hrs/timeperiod)	Wind	E_1^{ADP}	E_2^{ADP}	R_1^{ADP}	R_2^{ADP}
1	Constant	10	0	0.0008%	2.26%	0.02	0.006
2	Constant	10	✓	0.0003%	0.63%	0.04	0.007
3	Sinusoidal	10	0	0.0004%	0.02%	0.04	0.084
4	Sinusoidal	10	✓	0.0613%	2.67%	0.04	0.007
5	Fluctuating	10	0	0.0004%	0.04%	0.03	0.011
6	Fluctuating	10	✓	0.0000%	0.01%	0.03	0.005
7	Constant	1	0	0.0001%	0.06%	0.06	0.001
8	Constant	1	✓	0.0001%	0.08%	0.07	0.001
9	Sinusoidal	1	0	0.0012%	0.12%	0.03	0.002
10	Sinusoidal	1	✓	0.0020%	0.20%	0.03	0.001
11	Fluctuating	1	0	0.0000%	0.01%	0.02	0.001
12	Fluctuating	1	✓	0.0001%	0.02%	0.01	0.011

Table 1: Performance statistics of the ADP algorithm for different problem instances of a one-year energy model with the dispatch problem solved every H hours.

dynamic programming. In figure 7, we show the storage level produced by ADP for a single sample path, along with the storage level for each sample path when optimized knowing the entire sample path. The ADP solution has the intuitive behavior of building up storage earlier in the season to prepare for potentially lower rainfalls later in the season. By contrast, when we are allowed to optimize given the entire rainfall trajectory, the reservoir levels are held at lower levels until later in the season, reflecting the benefits of advance information. We believe that this illustrates the central weakness of models which sample uncertain information over the entire horizon, and then find the optimal solution given advance information.

7.4 Experiments on capacity acquisition

In the next set of experiments, we test the ability of the ADP algorithm to optimize capacity acquisition decisions over a multi-decade horizon under deterministic conditions. Given an initial set of capacities at the various conversion nodes, we allow the average demand to increase over the years. Eventually, the current level of capacities will be insufficient to cover the growing demands for energy. We now compare the decisions obtained using LP and ADP techniques to make capacity additions to the system. Figure 8 shows the objective function produced by the ADP algorithm over the optimal solution produced by the linear programming model as a function of the number of iterations of the ADP algorithm. The graph shows that we obtain a solution within 0.70% of the optimal after 200 iterations, and 0.24% of the optimal after 500 iterations.

Figure 7: Reservoir level produced by ADP (last iteration) for stochastic precipitation, and optimal reservoir levels produced by LP when optimized over individual scenarios.

Figure 8: The performance of the ADP algorithm as a function of the number of iterations.

Although the ADP algorithm produces an objective function close to the optimal objective function produced by the linear programming model, it is useful to compare the solution in terms of the investment decisions themselves. We found that the linear programming model reliably chooses to invest 100 percent in the lowest cost technology, even if there are competing technologies that have costs that are almost as low. By contrast, the ADP algorithm tends to spread its investment between technologies with comparable costs, just as we would expect with a stochastic model.

Figure 9: Comparison of investments in IGCC and pulverized coal using ADP and LP, at different levels of costs of either technology.

To demonstrate this behavior, we conducted a test involving several conversion technologies. In Figure 9, we illustrate the percentage investments in the lowest cost technology, namely pulverized coal, as a function of the cost differential between the two least expensive technologies (the second best technology being Integrated Gasification Combined Cycle (IGCC)). The unit capital costs and the unit hourly operating costs of these technologies are set such that IGCC is slightly more expensive to use than pulverized coal.

The linear programming model adds capacity exclusively to the pulverized coal technology at all four levels of cost. At the first three levels of cost differential, the ADP approach invests in both technologies in terms of capacity added, with the solutions approaching the linear programming solution as the cost difference increases. For the final level, where the cost difference is greatest, the ADP solution matches the linear programming solution. Despite the differences in the solutions for the lower levels of cost differential, we observed that the sum of the installed capacities of the two types is consistently similar for both the techniques. Moreover, the overall costs generated by the LP technique and the ADP algorithm were also found to match closely.

In table 2, we summarize the performance statistics after 500 iterations of the ADP algorithm for a number of problem scenarios. Here again, the ADP algorithm is seen to approach within one percent of the optimal solution quality for several different deterministic datasets.

Dataset	Demand type	Wind	Years	E_1^{ADP}	E_2^{ADP}	R_1^{cap}	R_2^{cap}
1	Constant	0	5	0.03%	0.85%	0.008	0.008
2	Constant	0	20	0.52%	1.37%	0.031	0.067
3	Constant	✓	20	0.24%	0.72%	0.027	0.061
4	Fluctuating	0	20	0.34%	0.92%	0.029	0.061
5	Fluctuating	✓	20	0.30%	0.91%	0.026	0.056

Table 2: Performance statistics of the ADP algorithm for different problem instances of a multi-year energy model.

7.5 Capacity additions with uncertainty

We close by reporting on a study involving uncertainty in a major policy decision. We test the model for an uncertain scenario where there could be a tax imposed on CO₂ emissions from the beginning of a certain year, T_{tax} , in the future. In the event of a CO₂ tax, a “correct” model would demonstrate a decrease in the investments in technologies that emit CO₂.

Our approach involved training the value function approximations for a certain number of iterations. In each iteration, the CO₂ tax is imposed in year T_{tax} with a certain probability in year 9. After a pre-determined number of training iterations, the value function approximations are kept fixed, and we then track investment decisions for each of the two scenarios. These decisions are shown in Figure 10. We observe that in the scenario where a carbon tax is imposed, the investments in carbon-based technologies decrease whereas the investments in renewable generation technologies increase. The algorithm exhibits a non-anticipative behavior, where the decisions up until the year T_{tax} are independent of future events.

8 Conclusions

We have presented a stochastic, multiscale energy resource planning model that handles competing energy resources and demands, a wide range of forms of uncertainty and the ability to model hourly variations over a multidecade horizon. The computational research shows that approximate dynamic programming can produce solutions that closely match the optimal solution of a deterministic linear programming model, while also providing solutions for problems that are far larger than what can be solved as a linear program.

The ADP model uses a proper model of the flow of information and decisions, and as a re-

Figure 10: Effect of uncertainty on investment decisions. Dashed lines represent decisions in the absence of a CO₂ tax.

sult avoids the short-cuts that have characterized many of the efforts at introducing uncertainty in optimization-based models. The experimental work shows that ADP produces robust behaviors such as storing rainfall earlier in the season than any of the solutions from a deterministic model acting on a specific scenario. It also exhibits nonanticipative behavior when faced with a major source of uncertainty such as the introduction of a carbon tax. While our algorithm is supported by some important theoretical results, we do not have a provably optimal algorithm for stochastic datasets which can serve as a comparable benchmark.

The major unresolved algorithmic issue involves the handling of coarse-grained uncertainty such as the imposition of a carbon tax or a breakthrough in technology. ADP has no difficulty simulating different scenarios describing these sources of uncertainty, but estimating a value function approximation that takes these issues into account in a proper way represents a new area of research. However, using suitably designed approximation strategies, ADP can produce good solutions that are nonanticipative (do not peek into the future) and which properly capture the dynamics in a realistic way.

The model can scale to handle spatially distributed energy production and consumption, but it is likely that additional work would be required to reduce execution times. When run on an aggregate dataset, total run times are on the order of several hours. If the problem were divided into 10 geographic regions, we suspect the run times would likely increase by a factor of 10, which is

manageable if unpleasant. It is also possible that training runs could be performed to estimate the value functions, after which testing runs could be performed (holding the value functions constant) for different policies. It is possible to perform these runs in parallel on a multiprocessor machine.

SMART was originally designed for the purpose of handling long-range policy studies, while retaining the accuracy of a fine-grained temporal model. However, it can also be used for studies addressing issues such as designing portfolios of energy generation technologies, and the choice of storage technologies that would complement a particular production portfolio. Thus, SMART can be used for a wide range of policy studies.

References

- Archibald, T. W., McKinnon, K. I. M. & Thomas, L. C. (1997), ‘An aggregate stochastic dynamic programming model of multi-reservoir systems’, *Water Resources Research* **33**(2), 333–340.
- Bertsekas, D. & Tsitsiklis, J. (1996), *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA.
- Birge, J. & Louveaux, F. (1997), *Introduction to Stochastic Programming*, Springer-Verlag, New York.
- Brown, P. D., Peas-Lopes, J. A. & Matos, M. A. (2008), ‘Optimization of pumped storage capacity in an isolated power system with large renewable penetration’, *IEEE Transactions on Power Systems* **23**(2), 523–531.
- Brunetto, C. & Tina, G. (2007), ‘Optimal hydrogen storage sizing for wind power plants in day ahead electricity market’, *IET Renewable Power Generations* **1**(4), 220–226.
- Castronuovo, E. D. & Lopes, J. A. P. (2004), ‘On the Optimization of the Daily Operation of a Wind-Hydro Power Plant’, *IEEE Transactions on Power Systems* **19**(3), 1599–1606.
- Cervellera, C., Chen, V. C. & Wen, A. (2005), ‘Optimization of a large-scale water reservoir network by stochastic dynamic programming with efficient state space discretization’, *European Journal of Operational Research* **171**, 1139–1151.
- De Ladurantaye, D., Gendreau, M. & Potvin, J.-Y. (2007), ‘Optimizing profits from hydroelectricity production’, *Computers and Operations Research* **00**, 00–00. doi:10.1016/j.cor.2007.10.012.
- Dupacova, J., Growe-Kuska, N. & Romisch, W. (2002), ‘Scenario reduction in stochastic programming’, *Mathematical programming* **95**(3), 493–512.
- Edmunds, T. A. & Bard, J. F. (1990a), ‘A decomposition technique for discrete time optimal control problems with and application to water resources management’, *Mathematical and Computer Modelling* **13**(1), 61–78.
- Edmunds, T. A. & Bard, J. F. (1990b), ‘Time-axis decomposition of large-scale optimal control problems’, *Journal of Optimization Theory and Applications* **67**(2), 259–277.
- Fan, J. & Gijbels, I. (1996), *Local Polynomial Modelling and Its Applications*, Chapman and Hall, London.
- Fishbone, L. & Abilock, H. (1981), ‘MARKAL - A linear programming model for energy systems analysis: Technical description of the BNL version’, *International Journal of Energy Research* **5**, 353–375.

- Foufoula-Georgiou, E. & Kitanidis, P. K. (1988), ‘Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems’, *Water Resources Research* **24**(8), 1345–1359.
- Garcia-Gonzalez, J., de la Muela, R. M. R., Santos, L. M. & Gonzalez, A. M. (2008), ‘Stochastic Joint Optimization of Wind Generation and Pumped-Storage Units in an Electricity Market’, *Power Systems, IEEE Transactions on* **23**(2), 460–468.
- Godfrey, G. & Powell, W. B. (2002), ‘An adaptive, dynamic programming algorithm for stochastic resource allocation problems I: Single period travel times’, *Transportation Science* **36**(1), 21–39.
- Hastie, T., Tibshirani, R. & Friedman, J. (2001), *The Elements of Statistical Learning*, Springer series in Statistics, New York, NY.
- Higle, J. & Sen, S. (1991), ‘Stochastic decomposition: An algorithm for two stage linear programs with recourse’, *Mathematics of Operations Research* **16**(3), 650–669.
- Higle, J. L. & Sen, S. (1996), *Stochastic Decomposition: A Statistical Method for Large Scale Stochastic Linear Programming*, Kluwer Academic Publishers.
- Hogan, W. (1975), ‘Energy policy models for project independence’, *Computers and Operations Research* **2**(2), 251–271.
- Høyland, K. & Wallace, S. (2001), ‘Generating scenario trees for multistage decision problems’, *Management Science* pp. 295–307.
- Jacobs, J., Freeman, G., Grygier, J., Morton, D., Schultz, G., Staschus, K. & Stedinger, J. (1995), ‘Socrates: A system for scheduling hydroelectric generation under uncertainty’, *Annals of Operations Research* **59**, 99–133.
- Johnson, S. A., Stedinger, J. R. & Staschus, K. (1991), ‘Heuristic operating policies for reservoir system simulation’, *Water Resources Research* **27**(5), 673–685.
- Johnson, T. L., DeCarolus, J. F., Shay, C. L., Loughlin, D. H., Gage, C. L. & Vijay, S. (2006), MARKAL Scenario Analyses of Technology Options for the Electric Sector, Technical report, United States Environmental Protection Agency.
- Kanudia, A. & Loulou, R. (1998), ‘Robust Responses to Climate Change via Stochastic MARKAL: the case of Québec’, *European Journal of Operations Research* **106**, 15–30.
- Kaut, M. & Wallace, S. (2003), ‘Evaluation of scenario-generation methods for stochastic programming’, *Stochastic Programming E-Print Series*.
- Korpaas, M., Holen, A. & Hildrum, R. (2004), ‘Operation and sizing of energy storage for wind power plants in a market system’, *International Journal of Electrical Power and Energy Systems* **25**(8), 599–606.
- Lamond, B. F. & Sobel, M. J. (1995), ‘Exact and approximate solutions of affine reservoir models’, *Operations Research* **43**(5), 771–780.
- Lamont, A. (1997), User’s Guide to the META*Net Economic Modeling System Version 2.0, Technical report, Lawrence Livermore National Laboratory.
- Lamont, A. D. (2008), ‘Assessing the long-term system value of intermittent electric generation technologies’, *Energy Economics* **30**(3), 1208–1231.
- Loulou, R., Goldstein, G. & Noble, K. (2004), Documentation for the MARKAL Family of Models, Technical report, Energy Technology Systems Analysis Programme.
- Messner, S. & Strubegger, M. (1995), User’s Guide for MESSAGE III, Technical report, International Institute for Applied Systems Analysis (IIASA).

- Messner, S., Golodnikov, A. & Gritsevskii, A. (1996), ‘A stochastic version of the dynamic linear programming model MESSAGE III’, *International Journal of Energy Research* **21**(9), 775–784.
- Nandalal, K. D. W. & Bogardi, J. J. (2007), *Dynamic Programming Based Operation of Reservoirs*, International Hydrology Series, Cambridge University Press.
- Nascimento, J. & Powell, W. B. (2009a), Optimal approximate dynamic programming algorithms for a general class of storage problems, Technical report, Department of Operations Research and Financial Engineering, Princeton University.
- Nascimento, J. M. & Powell, W. B. (2009b), ‘An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem’, *Mathematics of Operations Research* **34**(1), 210–237.
- NEMS (2003), The National Energy Modeling System: An Overview 2003, Technical report, Energy Information Administration, Office of Integrated Analysis and Forecasting, U.S. Department of Energy, Washington, DC 20585.
- Ormoneit, D. & Sen, S. (2002), ‘Kernel-based reinforcement learning’, *Machine Learning* **49**, 161–178.
- Paatero, J. V. & Lund, P. D. (2005), ‘Effect of energy storage on variations in wind power’, *Wind Engineering* **8**(4), 421–441.
- Pereira, M. & Pinto, L. (1991), ‘Multistage stochastic optimization applied to energy planning’, *Mathematical Programming* **52**, 359–375.
- Powell, W. B. (2007), *Approximate Dynamic Programming: Solving the curses of dimensionality*, John Wiley and Sons, New York.
- Powell, W. B. & Topaloglu, H. (2003), Stochastic programming in transportation and logistics, in A. Ruszczyński & A. Shapiro, eds, ‘*Handbooks in Operations Research and Management Science: Stochastic Programming*’, Vol. 10, Elsevier, Amsterdam, pp. 555–635.
- Powell, W. B. & Topaloglu, H. (2004), Fleet management, in S. Wallace & W. Ziemba, eds, ‘Applications of Stochastic Programming’, Math Programming Society - SIAM Series in Optimization, Philadelphia.
- Powell, W. B. & Topaloglu, H. (2005), Fleet management, in S. Wallace & W. Ziemba, eds, ‘Applications of Stochastic Programming’, Math Programming Society - SIAM Series in Optimization, Philadelphia, pp. 185–216.
- Powell, W. B., Ruszczyński, A. & Topaloglu, H. (2004), ‘Learning algorithms for separable approximations of stochastic optimization problems’, *Mathematics of Operations Research* **29**(4), 814–836.
- Powell, W. B., Shapiro, J. A. & Simão, H. P. (2002), ‘An adaptive dynamic programming algorithm for the heterogeneous resource allocation problem’, *Transportation Science* **36**(2), 231–249.
- Ruszczyński, A. (1993), ‘Parallel decomposition of multistage stochastic programming problems’, *Mathematical Programming* **58**(2), 201–228.
- Simao, H. P., Day, J., George, A. P., Gifford, T., Nienow, J. & Powell, W. B. (2009), ‘An approximate dynamic programming algorithm for large-scale fleet management: A case application’, *Transportation Science* **43**(2), 178–197.
- Sutton, R. & Barto, A. (1998), *Reinforcement Learning*, The MIT Press, Cambridge, Massachusetts.
- Swisher, J. R., Hyden, P. D., Jacobson, S. H., Schruben, L. W., Hosp, M. W. & Fredericksburg, V. A. (2000), ‘A survey of simulation optimization techniques and procedures’, *Simulation Conference Proceedings, 2000. Winter*.
- Topaloglu, H. & Powell, W. B. (2006), ‘Dynamic programming approximations for stochastic, time-staged integer multicommodity flow problems’, *Inform Journal on Computing* **18**(1), 31–42.

- Van Slyke, R. & Wets, R. (1969), ‘L-shaped linear programs with applications to optimal control and stochastic programming’, *SIAM Journal of Applied Mathematics* **17**(4), 638–663.
- Wallace, S. W. & Fleten, S. E. (2003), *Stochastic Programming*, Vol. 10, Elsevier Science, chapter Stochastic programming models in energy, pp. 00–00.

Appendix

The appendix provides a more detailed version of the model, followed by derivations of the derivatives needed to fit the value function approximations.

A.1 The full energy model

In this appendix, we develop a modeling framework that is able to capture different forms of uncertainty such as the future state of technology, climate changes and market information, as well as making decisions at various levels of time-discretization over a multi-decade horizon. We later test the framework on an aggregate model, largely to validate the algorithmic strategy. Our modeling framework, however, is designed to handle a fairly high level of detail, as well as a proper handling of uncertainty.

We proceed to describe the various elements of the model using a notational style developed specifically for stochastic, dynamic resource allocation problems. We then model two types of decisions in our model: the dispatch decisions to supply energy to the various sources of demand which will be performed on an hourly basis, and the capacity acquisition decisions to add new capacity to the conversion plants (or to build new plants for a newly discovered technology) which will be done every year. We note that dispatch decisions include decisions of how much to use, how much to store, and how much to withdraw from storage. We let $\mathcal{T}$ denote the set of once-a-year time-periods where we add capacity and $\mathcal{H}$ the set of hourly periods within a year, where $|\mathcal{H}| = H$ and $|\mathcal{T}| = T$.

In this section, we provide the notation needed to model all the elements of a dynamic system, which is organized along five fundamental dimensions: 1) system state variables, 2) decision variables, 3) exogenous information (random variables), 4) transition function and 5) objective function. In section 5, we use this modeling framework to formulate a deterministic linear programming model, and a stochastic optimization problem.

System state variables

We divide the state variable into the resource state (energy investments), storage (the amount of energy held in storage), and information about a range of system parameters.

Energy investments and storage

We let a represent a generic attribute vector. For example, the attribute vector of a conversion plant could be represented using

$$a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} \text{Type} \\ \text{Location} \\ \text{Age} \end{pmatrix}.$$

The attribute “Type” denotes the type of generation technology such as nuclear, gas turbine, photovoltaic or hydro-power. Age is typically in units of years. Location can be expressed at different levels of detail. For example, we may specify the location as a high-level region such as the central

valley in California or at the more detailed level of a zip code. We conduct our numerical experiments only for a spatially aggregate model, but our mathematical model is much more general than this, and the algorithmic strategy scales fairly easily to more complex problems.

It is useful to define the following attribute vector spaces:

- $\mathcal{A}$ = The set of all attribute vectors describing all possible investments in energy resources.
- $\mathcal{A}^{Res}$ = The set of all attribute vectors describing resource nodes, where raw energy resources enter the network (coal, oil, wind, etc.).
- $\mathcal{A}^{Conv}$ = The set of all attribute vectors describing conversion nodes, where raw resources are converted to electricity or other usable forms of energy.
- $\mathcal{A}^{Stor}$ = The set of all attribute vectors describing storage nodes.
- $\mathcal{A}^{Dem}$ = The set of all attribute vectors describing demand nodes.
- $\mathcal{A}'$ = The set of all attribute vectors describing all other types of nodes, which include collection and distribution nodes used to simplify the linear program,
 $= \mathcal{A} \setminus (\mathcal{A}^{Res} \cup \mathcal{A}^{Conv} \cup \mathcal{A}^{Stor} \cup \mathcal{A}^{Dem}).$

We define the connectivity of the energy network using

- $\vec{\mathcal{A}}(a)$ = The set of attribute vectors of nodes downstream in the dispatch network that are directly connected to a node with attribute vector a .
- $\overleftarrow{\mathcal{A}}(a)$ = The set of attribute vectors of nodes upstream in the dispatch network that are directly connected to a node with attribute vector a .

Using these attributes, the resource state variables can be described using

- R_{tha} = The total capacity in megawatts (MW) with attribute vector $a \in \mathcal{A}$ available in hour h of year t .
- R_{th} = The resource state vector in hour h of year t
 $= (R_{tha})_{a \in \mathcal{A}}.$
- y_{tha} = The amount of energy in storage for a node with attribute vector $a \in \mathcal{A}^{Stor}$ in hour h of year t .
- y_{th} = $(y_{tha})_{a \in \mathcal{A}^{Stor}}.$

Exogenously varying quantities

In our model, energy demand is treated as exogenous, although we recognize that an important dimension of the energy problem involves demand management. We model demands similarly to how we model energy resources. We let $a \in \mathcal{A}^{Dem}$ be the attributes of demand, which will include the type of energy demand (electricity, natural gas, oil for heating, gasoline for transportation) and could also include location (in a model which captures where demand is occurring). This allows us

to model demand as

$$\begin{aligned} D_{tha} &= \text{Demand with attribute } a \in \mathcal{A}^{Dem} \text{ at hour } h \text{ in year } t. \\ D_{th} &= (D_{tha})_{a \in \mathcal{A}}. \end{aligned}$$

Other exogenously varying parameters include

$$\begin{aligned} p_{th} &= \text{The amount of precipitation occurring in hour } h, \text{ year } t. \\ \rho_{th} &= \text{The information state in hour } h \text{ of time } t, \text{ which can include the state of technology } \\ &\quad (\rho_{th}^{Tech}), \text{ climate } (\rho_{th}^{Clim}), \text{ market } (\rho_{th}^{Mkt}) \text{ and exogenous energy supply } (\rho_{th}^{exo}). \\ &= (\rho_{th}^{Tech}, \rho_{th}^{Clim}, \rho_{th}^{Mkt}, \rho_{th}^{exo}). \end{aligned}$$

The vector $\rho_{th}^{exo} = (\rho_{tha}^{exo})_{a \in \mathcal{A}^{Conv}}$ is used to model intermittent energy at conversion facilities. If the resource a is nonintermittent (coal, nuclear, natural gas), then $\rho_{tha}^{exo} = 1$. If it is an intermittent supply such as wind or solar, we would have $\rho_{tha}^{exo} < 1$ whenever the exogenous supply of energy does not allow us to use the technology (such as a wind turbine) to its fullest capacity.

We finally define the system state variable using

$$\begin{aligned} S_{th} &= \text{The state of the system in hour } h \text{ of year } t \\ &= (R_{th}, y_{th}, D_{th}, p_{th}, \rho_{th}). \end{aligned}$$

The system state could include additional variables. For example, if there is a constraint imposed on the total CO₂ emitted since the beginning of the time horizon, then the state could have the added dimension of the current level of CO₂ emissions.

Decision variables

Here we describe the different types of decisions and the constraints they have to observe at a point in time.

$$\begin{aligned} x_{tha}^{cap} &= \text{The total capacity with attribute vector } a \in \mathcal{A}^{Conv} \text{ added to the system in hour } h \\ &\quad \text{of year } t. \\ x_{th}^{cap} &= (x_{tha}^{cap})_{a \in \mathcal{A}^{Conv}}. \\ x_{th,aa'}^{disp} &= \text{The flow from a node with attribute vector } a \text{ to a node with attribute vector } a' \\ &\quad \text{during hour } h \text{ of year } t. \\ x_{th}^{disp} &= (x_{th,aa'}^{disp})_{a \in \mathcal{A}, a' \in \mathcal{A}}. \end{aligned}$$

The vector of decisions is given compactly by

$$\begin{aligned} x_{th} &= \text{The vector of decisions in hour } h \text{ of year } t \\ &= (x_{th}^{cap}, x_{th}^{disp}). \end{aligned}$$

We defer to section 5 the presentation of deterministic and stochastic optimization models which will determine how the decisions are made.

We note here that, although we have modeled all decisions at the hourly level, real-life capacity addition decisions are seldom made more than once a year. This does not require us to change any of the notation that has already been presented, but we assume that for all intermediate hours $0 < h \leq H$, $x_{th}^{cap} = 0$. On the other hand, the dispatch decisions for distributing power to the various sources of demand have to be made every hour $h \in \mathcal{H}$.

These decisions are governed by a series of constraints that apply to decisions made at a point in time. For all conversion nodes, the output is constrained by the installed capacity at that node, given by

$$\sum_{a' \in \vec{\mathcal{A}}(a)} x_{th,aa'}^{disp} \leq \rho_{tha}^{exo} R_{tha}, \quad a \in \mathcal{A}^{Conv}. \quad (15)$$

Here, R_{tha} gives the installed capacity, while ρ_{tha}^{exo} captures the effect of intermittent sources such as wind or solar.

Conversion nodes are the primary source of constraint in the supply of energy. In our numerical work, we assume that supplies at the resource nodes (which captures the inputs of coal, natural gas, and so on) are unbounded, but otherwise we can write

$$\sum_{a' \in \vec{\mathcal{A}}(a)} x_{th,aa'}^{disp} \leq R_{tha}, \quad a \in \mathcal{A}^{Res}. \quad (16)$$

Here, the units of R_{tha} can be in barrels of oil, cubic feet of natural gas or tons of biomass. Unit conversions are handled in the flow conservation constraints, which are given by

$$\sum_{a' \in \vec{\mathcal{A}}(a)} \theta_{a'a} x_{th,a'a}^{disp} - \sum_{a'' \in \vec{\mathcal{A}}(a)} x_{th,aa''}^{disp} = 0, \quad a \in \mathcal{A} \setminus \mathcal{A}^{Stor}, \quad (17)$$

$$y_{tha} + \sum_{a' \in \vec{\mathcal{A}}(a)} \theta_{a'a} x_{th,a'a}^{disp} - \sum_{a'' \in \vec{\mathcal{A}}(a)} x_{th,aa''}^{disp} \geq 0, \quad a \in \mathcal{A}^{Stor}. \quad (18)$$

Equation (18) captures our ability to draw down from the amount that we have stored. $\theta_{a'a}$ captures changes in units (barrels of oil, cubic feet of gas and gallons of water to megawatts) as well as losses due to transmission and storage. We note that the sum over $a'' \in \vec{\mathcal{A}}(a)$ includes flows from storage, and the sum over $a' \in \vec{\mathcal{A}}(a)$ includes flows that remain in storage. θ_{aa} , for $a \in \mathcal{A}^{stor}$, would capture losses of energy held in storage. We note that equations (17)-(18) represent constraints at a point in time. Later we present the equations that govern how storage levels evolve over time.

For all demand nodes, the total input should match the total demand,

$$\sum_{a' \in \vec{\mathcal{A}}(a)} x_{th,a'a}^{disp} = \hat{D}_{tha}, \quad \text{where } a \in \mathcal{A}^{Dem}, \quad (19)$$

where we assume that there is some expensive, exogenous source (such as imported electricity) to handle situations when all other sources are not sufficient to meet demand. Finally, we require

$$x_{th}^{disp}, x_{th}^{cap} \geq 0. \quad (20)$$

Let $\mathcal{X}_{th}^{disp}$ be the feasible region defined by equations (15) - (20).

Exogenous information

Exogenous information is expressed in the form of random variables governed by probability distributions or an exogenously provided file of scenarios. All random variables are denoted by “hats” (as in $\hat{D}$).

- $\hat{R}_{th}$ = Exogenous changes in the resource state in hour h of year t (due, for example, to equipment failures or weather damage).
- $\hat{D}_{tha}$ = Change in demand with attribute vector $a \in \mathcal{A}^{Dem}$ in hour h of year t .
- $\hat{\rho}_{th}$ = Exogenous changes in any parameters governing technology, markets, climate and exogenous energy sources,
 $= (\hat{\rho}_{th}^{Tech}, \hat{\rho}_{th}^{Mkt}, \hat{\rho}_{th}^{Clim}, \hat{\rho}_{th}^{exo})$.
- $\hat{p}_{th}$ = The change in precipitation in hour h of year t .

We summarize all random variables in a single variable W_t (without a hat), as follows:

$$\begin{aligned} W_{th} &= \text{A vector of all exogenous information regarding parameters appearing in hour } h \text{ of} \\ &\quad \text{year } t \text{ which can include data on weather, demands, technological parameters and} \\ &\quad \text{market information,} \\ &= (\hat{R}_{th}, \hat{D}_{th}, \hat{\rho}_{th}, \hat{p}_{th}). \end{aligned}$$

The availability of technologies of each type would be an element of $\hat{\rho}_{th}^{Tech}$. Examples of components of $\hat{\rho}_{th}^{Clim}$ might be the level of allowable emissions out to year T , and the total quantities and hourly patterns of wind (which is a specific example of the exogenous production factor data, ρ_{tha}^{exo}). Interest rates, exogenous resource prices, and resource availability, would be included in $\hat{\rho}_{th}^{Mkt}$. Although all of these have been indexed at the hourly level, we note that each of these could be arriving on different time-scales. For example, information regarding climate data might be available to be used by the system at anywhere from an hourly to monthly frequency. Prices might change on a daily basis, while technological developments could take as long as a year to multiple years to impact the system.

Later, we need a more formal vocabulary to describe the stochastic process driving our system. Let Ω be the set of all possible realizations of the sequence $(W_{11}, \dots, W_{1H}, W_{21}, \dots, W_{2H}, \dots, W_{TH})$ where T is the number of years and H is the number of time periods in a year. We let $W_{th}(\omega)$ be a sample realization of the random variables in W_{th} when we are following sample path ω . To complete our formalism, let $\mathcal{F}$ be the sigma-algebra on Ω , with filtrations $\mathcal{F}_{th} \subset \mathcal{F}_{t,h+1}$ and $\mathcal{F}_{tH} \subset \mathcal{F}_{t+1,0}$,

where $\mathcal{F}_{th}$ is the sigma-algebra generated by the information up through year t , hour h . We use this notation below to describe our algorithm and establish properties of the decision function.

Transition function

The state of the system evolves as a result of the decisions taken as well as random events, which is written as,

$$\begin{aligned} S_{t,h+1} &= S^M(S_{th}, x_{th}, W_{t,h+1}), \quad h = 0, 1, \dots, H-1, \\ S_{t+1,0} &= S^M(S_{tH}, x_{tH}, W_{t+1,0}). \end{aligned}$$

Elements of the system state may transition in different ways. For example, new capacity may be added to the system and old ones may be retired, giving us

$$R_{t,h+1,a} = R_{tha} + x_{tha}^{cap} + \hat{R}_{t,h+1,a}, \quad a \in \mathcal{A}^{Conv}, \quad (21)$$

where x_{tha}^{cap} is positive for capacity additions and negative when capacities are retired. $\hat{R}_{t,h+1,a}$ involves changes in the resource state as a result of unforeseen events such as Hurricane Katrina.

Changes in demand, precipitation and the parameter vector ρ_t are all handled similarly using

$$\hat{D}_{t,h+1} = D_{th} + \hat{D}_{t,h+1}, \quad (22)$$

$$\hat{p}_{t,h+1} = p_{th} + \hat{p}_{t,h+1}, \quad (23)$$

$$\rho_{t,h+1} = \rho_{th} + \hat{\rho}_{t,h+1}. \quad (24)$$

Let y_{tha} be energy stored in hour h of year t associated for storage location $a \in \mathcal{A}^{Stor}$. In our numerical work, we model the water reservoir as the only form of storage, but our model is much more general. We can represent the transition equation for storage facilities using

$$y_{t,h+1,a} = y_{tha} + \sum_{a' \in \vec{\mathcal{A}}(a)} \theta_{a'a} x_{th,a'a}^{disp} - \sum_{a'' \in \vec{\mathcal{A}}(a)} x_{th,aa''}^{disp} + \hat{p}_{t,h+1,a}, \quad a \in \mathcal{A}^{Stor} \quad (25)$$

where $\hat{p}_{t,h+1,a}$ captures exogenous precipitation into the reservoir with attribute a . A slightly modified equation is used when $h = H$. We note that the only flow between time periods occurs through storage nodes. All other flows of energy occur within a time period.

Cost Functions

The costs consist of the costs for adding new capacity, the costs related to energy dispatch which cover the physical movement of energy resources and transmission of electrical energy, and costs for

purchasing fuel from the various energy sources.

- c_{tha}^{cap} = The capital cost for adding a unit of capacity with attribute vector $a \in \mathcal{A}^{Conv}$ in hour h of year t .
- $c_{th,aa'}^{disp}$ = The cost of unit flow from node with attribute vector $a \in \mathcal{A}$ to node with attribute vector $a' \in \mathcal{A}$ in hour h of year t (for example, the transportation cost of moving biomass).
- c_{tha}^f = Unit cost of fuel corresponding to resource node with attribute vector $a \in \mathcal{A}^{Res}$ in hour h of year t .
- c_{tha}^p = Unit operating cost of a conversion plant with attribute vector $a \in \mathcal{A}^{Conv}$ in hour h of year t (assumed to be proportional to the flow).

Total capacity and dispatch costs are given by

$$\begin{aligned}
C_{th}^{cap}(S_{th}, x_{th}^{cap}) &= \text{The incremental capital costs in hour } h \text{ of year } t, \\
&= \sum_{a \in \mathcal{A}} c_{tha}^{cap} x_{tha}^{cap}. \\
C_{th}^{disp}(S_{th}, x_{th}^{disp}) &= \text{The total costs resulting from the dispatch of energy in hour } h \text{ of year } t, \\
&= \sum_{a \in \mathcal{A}} \sum_{a' \in \mathcal{A}} c_{th,aa'}^{disp} x_{th,aa'}^{disp} + \sum_{a \in \mathcal{A}^{Res}} \sum_{a' \in \mathcal{A}} c_{tha}^f x_{th,aa'}^{disp} + \sum_{a \in \mathcal{A}^{Conv}} \sum_{a' \in \mathcal{A}} c_{tha}^p x_{th,aa'}^{disp}.
\end{aligned}$$

We may express the total cost function in hour h of year t using,

$$C_{th}(S_{th}, x_{th}) = C_{th}^{cap}(S_{th}, x_{th}^{cap}) + C_{th}^{disp}(S_{th}, x_{th}^{disp}).$$

A.2 Calculating the derivative $\hat{v}^{hydro}$

We begin by computing the marginal value of additional energy in a single, scalar storage which, for our numerical experiments, was represented as a water reservoir. We do this by computing the marginal impact of having an additional unit of water in storage at a point in time. This calculation is then used in the derivation of the derivative for additional capacity, which impacts an entire year.

The objective function at hour h is given by

$$F_{th}^\pi(S_{th}) = -C_{th}(S_{th}, X_{th}^\pi) - \sum_{h' > h} C_{th'}(S_{th'}, X_{th'}^\pi) - \sum_{t' > t} \sum_{h \in \mathcal{H}} C_{t'h}(S_{t'h}, X_{t'h}^\pi) \quad (26)$$

$$\begin{aligned}
&= -C_{th}(S_{th}, X_{th}^\pi) + V_{th}^x(S_{th}^x), \\
&\approx -C_{th}^{disp}(S_{th}, X_{th}^\pi) - C_{th}^{cap}(S_{th}, X_{th}^\pi) + \bar{V}_{th}^{hydro}(y_{th}^x) + \bar{V}_{th}^{cap}(R_{th}^x).
\end{aligned} \quad (27)$$

Note that $C_{th}^{cap}(S_{th}, X_{th}^\pi) = 0$ for all $h > 0$ since we assume that capacity addition decisions are made only at the beginning of each year.

To update $\{\bar{V}_{th}^{hydro}(y_{th}^x)\}_{h \in \mathcal{H}, t \in \mathcal{T}}$, we need to differentiate $F_{th}^\pi(S_{th})$ with respect to y_{th} which is given by

$$\hat{v}_{th}^{hydro} = \frac{\partial F_{th}^\pi(S_{th})}{\partial y_{th}}.$$

We note that we differentiate around the pre-decision state y_{th} , but then use this to update $\bar{V}_{t,h-1}^x(y_{t,h-1}^x)$ around the previous post-decision state $y_{t,h-1}^x$. We may reasonably claim that

$$\frac{\partial R_{th}^x}{\partial y_{th}} \approx 0 \quad (28)$$

which captures the belief that an incremental change in the amount of energy in storage will have virtually no impact on our investments in energy resources. For example, an incremental increase in the amount stored in a reservoir will not have a measurable impact on the number of wind turbines we would purchase next year. This approximation gives us

$$\hat{v}_{th}^{hydro} = \frac{\partial F_{th}^\pi(S_{th})}{\partial y_{th}} \quad (29)$$

$$\begin{aligned} &\approx \frac{\partial}{\partial y_{th}} \left(-C_{th}^{disp}(S_{th}, X_{th}^\pi) + \bar{V}_{th}^{hydro}(y_{th}^x) + \bar{V}_{th}^{cap}(R_{th}^x) \right) \\ &= \frac{\partial}{\partial y_{th}} \left(-C_{th}^{disp}(S_{th}, X_{th}^\pi) + \bar{V}_{th}^{hydro}(y_{th}^x) \right) + \frac{\partial \bar{V}_{th}^{cap}(R_{th}^x)}{\partial R_{th}^x} \frac{\partial R_{th}^x}{\partial y_{th}} \\ &\approx \frac{\partial}{\partial y_{th}} \left(-C_{th}^{disp}(S_{th}, X_{th}^\pi) + \bar{V}_{th}^{hydro}(y_{th}^x) \right) \end{aligned} \quad (30)$$

where we use the approximation, $\partial R_{th}^x / \partial y_{th} \approx 0$. The derivative (30) is approximated using a finite difference. Let S_{th}^+ be the resource state where y_{th} is incremented by 1.

We then approximate the derivative using

$$\hat{v}_{th}^{hydro} \approx F_{th}^{disp,\pi}(S_{th}^+) - F_{th}^{disp,\pi}(S_{th}),$$

where $F_{th}^{disp,\pi}(S_{th}) = -C_{th}^{disp}(S_{th}, X_{th}^\pi) + \bar{V}_{th}^{hydro}(y_{th}^x)$, denotes the dispatch optimization component of the objective function. This numerical derivative is exceptionally fast to compute since it involves a trivial reoptimization of a small linear program.

It is important to emphasize that the calculation of $\hat{v}_{th}^{hydro}$ captures not only the hourly variability of wind, but also the uncertainty, since the value of additional stored energy does not require knowing the future about wind. The marginal value is also computed in the context of a storage facility.

A.3 Calculating the derivative $\hat{v}^{cap}$

The procedure for updating $\{\bar{V}_{th}^{cap}(R_{th}^x)\}_{h \in \mathcal{H}, t \in \mathcal{T}}$ is somewhat more involved. We assume that investment decisions are made once each year, taking effect in a future year. Our goal is to calculate

the marginal value of an additional unit of energy conversion capacity from a decision made in a previous year.

Changing the resource vector at the beginning of the year changes the resource vector for the entire year through equation (15). Also, since we only make capacity decisions once each year, we do not compute $\bar{V}_{th}^{cap}(R_{th}^x)$ for all h , but rather only for $h = 0$. For this reason, we have to capture the marginal value of additional capacity on both the yearly investment problem (11) as well as all the hourly dispatch problems.

We start with F_{t0}^π given by equation (26), although we will use the approximation in (27). Our goal is to compute

$$\hat{v}_{t0a}^{cap} = \frac{\partial F_{t0}^\pi(S_{t0})}{\partial R_{t0a}}. \quad (31)$$

F_{t0}^π can be written as a function of $R_{t0}, \dots, R_{th}, \dots, R_{t,H-1}$, recognizing that $R_{th} = R_{t0}$ for $0 \leq h < H$. R_{th} impacts the dispatch decision at hour h , as well as the post-decision storage level y_{th}^x which impacts future problems. In addition, changing R_{t0} has an impact on $R_{t+1,0}$, which impacts decisions in future years. To capture these interactions, we write

$$\frac{dF_{t0}^\pi(S_{t0})}{dR_{t0a}} = \left(-\frac{\partial C_{t0}^{cap}(S_{t0}, X_{t0}^\pi)}{\partial R_{t0a}} + \frac{\partial F_{t+1,0}^\pi}{\partial R_{t+1,0}} \frac{\partial R_{t+1,0}}{\partial R_{t0}} \right) - \sum_{0 \leq h < H} \frac{\partial F_{t0}^\pi}{\partial R_{th}} \frac{\partial R_{th}}{\partial R_{t0}}. \quad (32)$$

We can compute the marginal value of additional resources for the capacity subproblem at hour $h = 0$, given by equation (11), by using either the dual variable for equation (12) or a finite difference. We redefine S_{t0}^+ here to be S_{t0} with R_{t0} incremented by 1. The finite difference is then given by

$$\hat{v}_{t0a}^+ = F_{t0}^{cap,\pi}(S_{t0}^+) - F_{t0}^{cap,\pi}(S_{t0}).$$

This approximates the two terms in the set of parentheses on the right hand side of (32).

We next observe that

$$\frac{\partial F_{t0}^\pi}{\partial R_{th}} = \frac{\partial F_{th}^\pi}{\partial R_{th}}$$

since changing R_{th} has no impact on the problem for earlier hours, and

$$\frac{\partial R_{th}}{\partial R_{t0}} = 1.$$

The derivative $\partial F_{th}^\pi / \partial R_{th}$ involves finding the derivative of $F_{th}^\pi(S_{th})$ approximated by (27), which we compute using the dual variable for equation (15). Let ν_{tha}^{disp} be the dual variable for (15). Since the right hand side of (15) is $\rho_{tha}^{exo} R_{tha}$, the marginal impact of increasing R_{tha} on the dispatch optimization component of the objective function is given by

$$\frac{\partial F_{th}^\pi(S_{th})}{\partial R_{tha}} = \frac{\nu_{tha}^{disp}}{\rho_{tha}^{exo}}.$$

We finally pull these calculations together to find the marginal value of an additional unit of energy resource at the beginning of the year using

$$\hat{v}_{t0a}^{cap} \approx \sum_{h \in \mathcal{H}} \frac{\nu_{tha}^{disp}}{\rho_{tha}^{exo}} + \hat{v}_{t0a}^+. \quad (33)$$

We use $\hat{v}_{t0a}^{cap}$ to update the value function approximation $\bar{V}_{t-1,0a}^{cap}$.