



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Unit Commitment Without Commitment: A Dynamic Programming Approach for Managing an Integrated Energy System Under Uncertainty

David B. Brown, James E. Smith

To cite this article:

David B. Brown, James E. Smith (2025) Unit Commitment Without Commitment: A Dynamic Programming Approach for Managing an Integrated Energy System Under Uncertainty. *Operations Research* 73(4):1744-1766.
<https://doi.org/10.1287/opre.2023.0546>

Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Operations Research*. Copyright © 2025 The Author(s). <https://doi.org/10.1287/opre.2023.0546>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2025 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Contextual Areas

Unit Commitment Without Commitment: A Dynamic Programming Approach for Managing an Integrated Energy System Under Uncertainty

David B. Brown,^{a,*} James E. Smith^b^aFuqua School of Business, Duke University, Durham, North Carolina 27708; ^bTuck School of Business, Dartmouth College, Hanover, New Hampshire 03755

*Corresponding author

Contact: dbbrown@duke.edu,  <https://orcid.org/0000-0002-5458-9098> (DBB); jim.smith@dartmouth.edu, <https://orcid.org/0000-0002-9429-7567> (JES)

Received: October 9, 2023


Revised: May 30, 2024; August 11, 2024;
November 22, 2024

Accepted: December 3, 2024

Published Online in Articles in Advance:
March 5, 2025Area of Review: Environment, Energy, and
Sustainability<https://doi.org/10.1287/opre.2023.0546>

Copyright: © 2025 The Author(s)

Abstract. Though variability and uncertainty have always posed challenges for power systems, the increasing use of renewable energy sources has exacerbated these issues. At a vertically integrated utility, the system operator manages many generation units—renewable and otherwise—and storage units to ensure that the total energy produced matches contemporaneous demand. Current industry practice at these utilities involves solving “unit commitment” and “economic dispatch” optimization problems to choose production plans. These models, though complex, do not explicitly incorporate uncertainty. In this paper, we develop a dynamic programming approach to help system operators manage production under uncertainty. We formulate the problem as a stochastic dynamic program and use Lagrangian methods to decompose the system across units. The Lagrangian model relaxes the demand-matching constraint and introduces stochastic Lagrange multipliers that can be interpreted as prices representing the varying marginal value of energy production; each unit is then operated to maximize its own expected “profit” given these uncertain prices. These unit-specific value functions are then used to incorporate longer-term effects in dispatch decisions. The unit-specific value functions also provide a way to value generation and storage units in an uncertain environment. We develop relevant theory and demonstrate this dynamic approach using data from the Duke Energy Carolinas and Progress systems. Our numerical experiments demonstrate that this dynamic approach is computationally feasible at an industrial scale and can improve current practice. Specifically, our results suggest that this dynamic approach can reduce operational costs by about 2% on average in the present Duke Energy system and, in a “future” system with increased solar and storage capacity, can reduce operational costs by 4%–5% on average. Perhaps more strikingly, this dynamic approach, on average, performs within 0.2%–0.3% of production plans based on perfect foresight about future net demands.

 **Open Access Statement:** Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Operations Research*. Copyright © 2025 The Author(s). <https://doi.org/10.1287/opre.2023.0546>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This work was supported by the Department of Energy Advanced Research Projects Agency - Energy [Grant DE-AR0001283, “A Grid that’s Risk-Aware for Clean Electricity (GRACE)”] and Fuqua and Tuck.

Supplemental Material: The electronic companion is available at <https://doi.org/10.1287/opre.2023.0546>.


Keywords: energy • weakly coupled stochastic dynamic programs • unit commitment problems

1. Introduction

Though variability and uncertainty have always posed challenges for power systems, the growing use of renewable energy has exacerbated these issues. To illustrate, consider the “duck curve” of Figure 1 showing a snapshot of net load on spring days in California, as the solar energy supply has grown over the last few years. As shown there, the increased reliance on solar energy

forces operators to quickly ramp up production when the sun sets and solar production falls. In a setting where unpredictable thunderstorms sometimes “pop up” as the sun sets (for example, in the southeastern United States), such storms can suddenly reduce the supply of solar energy and require the system operator to ramp up production from other sources quickly. Some generation units—for example, thermal units—

Figure 1. (Color online) California ISO's "Duck Curve"



Source. U.S. Energy Information Administration (October 2023).

may have significant startup costs and limited ability to ramp up production to meet such unexpected increases in demand. Other generation units, such as gas turbines, may be more expensive to operate but can ramp up or down quickly. Managing uncertainty cost-effectively requires using an appropriate mix of such technologies.

In a vertically integrated energy system (where the system operator controls the operation of all plants in the system), systems are typically managed by solving "unit commitment" (UC) and "economic dispatch" (ED) problems (see, e.g., Conejo and Baringo 2018). Each morning, the system operator solves a UC problem to determine the "commitment" of generating units to minimize operating costs while meeting the expected demand over the course of the day or week. Then, once the commitment of generating units is decided, each hour (or more frequently), the system operator solves the ED problem to determine the actual power output of each generating unit to meet the realized demands at minimum cost, subject to the commitments determined in the UC problem as well as other constraints of the system. In particular, in the ED problem, slow-starting units (such as thermal plants) are "committed" to be on or off according to the result of the UC problem, though their production levels may be adjusted; that is, the binary variables corresponding to on/off for the slow-start units in the ED problem are set to their optimal values in the UC problem. The flows into and out of pumped-storage hydro are similarly committed to the UC solution. If additional power is needed, the ED problem can dispatch fast-starting units (e.g., gas turbines), subject to their physical constraints. These ED problems are typically myopic, focusing on minimizing the costs for a specific period, given the current state of the system. These models, though complex, typically do not explicitly consider uncertainty in supply or demand.

In this paper, we describe a dynamic approach for managing an integrated energy system under uncertainty based on methods from weakly coupled

stochastic dynamic programming (e.g., Hawkins 2003, Adelman and Mersereau 2008). Specifically, we model demand and renewable supply (and potentially other uncertainties)—the *world-state*—as an exogenous discrete-time stochastic process over a fixed horizon, say, a day or a week. Periods in the model correspond to hours, and dispatch decisions are made in each period. The modeled power system includes many generation units, each with different characteristics, as well as storage units. We describe the system model at a high level in Section 2 and the generation and storage unit models in Section 4.

Although the system-level stochastic dynamic program (DP) is too complex to solve exactly, we can decompose the problem using a Lagrangian relaxation. In particular, as discussed in Section 3, we relax the constraints that demand and production must balance in each period and each scenario and impose Lagrange multipliers that punish violations of these constraints. These Lagrange multipliers are, in general, *stochastic*—depending on the history of world-states—and can be interpreted as "prices" that units are paid for the energy produced. The resulting relaxed model decouples across units into a set of unit-specific DPs where each unit maximizes its own profit, keeping track of its own state and the stochastic world-state.

We can consider various functional forms for the stochastic price models in this Lagrangian relaxation. For any price model, the relaxed model provides an upper bound on the system's total profit (or a lower bound on the total costs), and we can find the best such bound for a particular price model by solving the dual optimization problem. The optimal Lagrange multipliers (or prices) ensure that the production in the relaxed model matches certain statistical features of the demand process. For example, with a fully general stochastic price process, the optimality conditions ensure that production in the Lagrangian model matches demand in every scenario, albeit with mixed policies. In this case, we can

provide a bound on the gap between the Lagrangian and the value function that is independent of the number units in the system and the number of world-states in the model. If we consider period-specific deterministic prices, the optimality conditions ensure production in the Lagrangian model matches demand “in expectation” in each period.

In Section 4, we describe specific models for the production and storage units and derive structural properties of these unit-specific DPs. These structural properties greatly simplify the solution of the unit-specific DPs and help make the decomposed Lagrangian model tractable.

We see this decomposed Lagrangian DP as serving a role analogous to that of the deterministic UC problems in current practice: the relaxed DP model provides a plan for operating the system on a given day. Here, however, the plans are state contingent, describing what each unit should do in each world-state. To ensure that these unit-specific state-contingent plans are consistent and meet the actual realized demand in each period, we solve a *forward-looking* version of the ED problem that is described in Section 5. Specifically, in each period, we solve a mixed-integer linear program (MILP) that maximizes the sum of unit-specific values (using unit-specific DP value functions from the Lagrangian relaxation) subject to the constraint of matching demand exactly and respecting all other system constraints. These unit-specific value functions thus embed long-term considerations when making hourly dispatch decisions. In this forward-looking ED problem, we assume the system operator has full flexibility to control any and all plants as well as storage; that is, it operates *without commitment*. The ramping constraints of slow-starting units (as well as all other units) are fully respected in this model, but in contrast to current practice, in this dynamic approach, we assume that there are no exogenous constraints on the ED problem imposed by the solution of the UC problem. Thus, in our forward-looking version of the ED problem, the system operator can start up or shut down slow-start units (as well as fast-start units) and deploy storage as desired.

In Section 6, we evaluate the performance of this proposed dynamic approach using Monte Carlo simulation: we consider various price models in this dynamic approach and compare performance to current practice as well as performance bounds provided by a perfect information relaxation (Brown et al. 2010). The simulations are set in the context of the Duke Energy Carolinas and Duke Energy Progress system (serving North and South Carolina), and the results suggest that the proposed dynamic approach is feasible for large-scale problems and performs quite well. In Electronic Companion Section EC-5, we provide pseudocode for implementing the main elements of our proposed approach and results for a fully specified example.

The work described in this paper was undertaken as part of the larger GRACE project (a Grid that is Risk

Aware for Clean Electricity, funded by the U.S. Department of Energy) that involves researchers from several universities and Duke Energy as an industry partner (Patino-Echeverri 2023). Our numerical experiments use confidential data about Duke Energy’s system provided as part of this project. In these experiments, we also use demand-forecasting models that were developed by other GRACE team members and are based on confidential Duke Energy data.

1.1. Literature Review

Tahanan et al. (2015) provide a comprehensive review of work on unit commitment under uncertainty. In practice, operators frequently tweak their deterministic UC models in response to uncertainty. For example, a common approach for managing uncertainty in supply/demand is for operators to augment their deterministic UC models to include reserve constraints that ensure capacity will be available if demand is higher than expected or if the system were to experience a sudden loss of supply. Operators also sometimes rerun their deterministic UC models if conditions deviate significantly from their forecasts or adjust dispatch plans in ad hoc ways, particularly plans related to storage.¹

Tahanan et al. (2015) distinguish three approaches to the UC problem that explicitly consider uncertainty: stochastic optimization, robust optimization, and chance-constrained optimization. We use a stochastic optimization approach and will concentrate our literature review on that area. Most papers in this literature rely on scenario tree models of uncertainty. One can solve the resulting formulations using mixed-integer programming techniques, but the problem size explodes as the number of scenarios grows. Tahanan et al. (2015) distinguish three decomposition approaches to deal with this challenge: scenario decomposition, Benders decomposition, and unit decomposition.

Scenario decomposition involves relaxing the nonanticipativity constraints at each node of the scenario tree and optimizing corresponding dual variables. Takriti et al. (1996, p. 1503) develop this approach using progressive hedging techniques and note “the execution time of the algorithm grows rapidly as the number of scenarios included increases and their demands are more diverse.” Takriti and Birge (2000) show that the duality gap for this approach is bounded from above by a term that grows linearly with the number of branching points in the scenario tree. Benders decomposition has also been applied widely, typically in two-stage models of demand uncertainty. Cerisola et al. (2009) demonstrate Benders decomposition approaches in multistage power production planning on numerical examples involving eight scenarios.

Our approach is a form of unit decomposition that uses Lagrangian relaxation of the constraints that couple the units across the system. Several researchers (e.g., Carpentier et al. 1996, Dentcheva and Römisich 1998,

Nowak and Römisich 2000) develop unit decomposition approaches with stochastic Lagrange multipliers in power planning problems with scenario trees. Carpentier et al. (1996, p. 1068) note that “scenario trees have to be kept simple enough to avoid combinatorial explosion of computation.” In this stream of the unit decomposition literature, solving the unit-specific subproblems using dynamic programming is standard, with the size of these subproblems growing exponentially with the number of scenarios (see, e.g., section 3.2 of Nürnberg and Römisich 2002). In our approach, we work with simpler forms of stochastic Lagrange multiplier (or price) models to avoid this combinatorial explosion.

Our work also falls within and contributes to the literature on approximate dynamic programming techniques in energy planning applications. Papers in this stream model uncertainties using general Markov processes rather than scenario trees. For example, Khazaei and Powell (2018) develop a lookahead-based policy using parameter tuning for managing daily production operations within a model that optimizes long-term investments in renewable energy sources. Powell and Meisel (2015) provide a tutorial on using tuned parametric approximations of value functions and lookahead techniques, focusing on energy storage applications. Barty et al. (2010a) develop a stochastic Lagrangian decomposition approach similar to ours using a weakly coupled stochastic DP; they illustrate their approach with autoregressive (AR) Lagrange multipliers on a simple power planning problem involving two hydraulic plants and one thermal generation unit. Barty et al. (2010a) note that a key advantage of the approach is that it scales linearly in the number of units. Barty et al. (2010b) further develop this approach and illustrate it in an example with seven hydraulic plants and one thermal generation unit.

Finally, our paper contributes to the growing body of methodological research on weakly coupled stochastic DPs (Hawkins 2003, Adelman and Mersereau 2008), which generalizes ideas from the restless bandit literature (e.g., Whittle 1988). Researchers have successfully applied weakly coupled DP methods to a variety of problems, including marketing (Bertsimas and Mersereau 2007), assortment planning (Caro and Gallien 2007, Brown and Smith 2020), network revenue management (Topaloglu 2009), and multilocation inventory management (Brown and Zhang 2022, Miao et al. 2022), among many others. Brown and Smith (2020) and Brown and Zhang (2022) provide results that show Lagrangian-based heuristics are asymptotically optimal as the systems grow large (e.g., as the number of units grows large). Though these asymptotic optimality results are for different settings, these results and earlier results for deterministic systems (e.g., Bertsekas 1996, proposition 5.26) suggest that Lagrangian-relaxation methods may perform well in this setting. As mentioned above, we

provide a duality gap result for the Lagrangian relaxation (see Theorem 2), albeit for a complete price model that we cannot solve in complex examples. In our numerical examples in Section 6, we use perfect information performance bounds to show that our methods yield costs that are within 0.2%–0.4% of optimal in these examples.

1.2. Contributions

We view our contributions to be both applied and methodological. On the applied front, our approach is a scalable and practically viable method for managing production in an integrated energy system facing uncertainty. Our numerical experiments based on the Duke Energy system demonstrate that our approach performs well at an industrial scale, taking seconds to minutes to solve the Lagrangian dual problems on a desktop computer, outperforming current practice in realistic settings and achieving within a fraction of a percent of optimal performance given perfect information about future demands. On the methodological front, we advance the analysis and application of stochastic Lagrange multipliers in weakly coupled stochastic DPs and develop new structural properties for the resulting unit-specific value functions that simplify their solution.

2. The System Model

We consider a finite horizon with periods $t = 1, \dots, T$ with power demands (d_1, \dots, d_T) in each period. In our numerical examples, we will consider periods of one hour in duration and a time horizon $T = 24$ or 168 corresponding to a day or a week. The system operator or decision maker (DM) seeks to maximize the profit (or minimize costs) over this horizon.

For now, we consider a general model of demands and renewable supplies, assuming they are generated by some Markov process with world-state $\psi_t \in \Psi$. We assume that the world-state includes the current demand d_t and supplies for weather-dependent units and everything needed to generate forecasts for future demands and supplies for weather-dependent units. This world-state ψ_t could, in principle, be a large and complex state variable noting current temperatures and forecasts of future temperatures (as well as other weather variables), plant outages, and fuel costs, as well as the current demand and weather-related supplies. We assume that the world-state transitions are exogenous (independent of the system state and DM’s actions) and that the period- t world-state ψ_t is known to the DM when making decisions in period t . We let $\mathbb{F} = (\mathcal{F}_1, \dots, \mathcal{F}_T)$ denote the filtration representing the DM’s knowledge of the world-state over time. To avoid measurability and related technical issues, we will assume that the world-state space Ψ is finite.

There are S generation units of various types, including both production and storage units, with details to be described in Section 4. The state of unit s in any given period is summarized by a state variable $x_s \in X_s$. In each period, the DM selects an action a_s from a feasible set $A_s(x_s, \psi_t) \subseteq A_s$, where A_s denotes the action space. These actions produce $p_s(a_s)$ units of energy and cost $c_s(x_s, a_s, \psi_t)$. The unit's state then evolves deterministically to $\chi_{s,t}(x_s, a_s, \psi_t)$ in the next period. The constraint sets may depend on the world-state ψ_t , reflecting the availability of wind or solar power or a unit or system outage. Similarly, the costs may depend on ψ_t , reflecting, for example, fuel costs for generation units. The transitions may also depend on the world-state, reflecting, for example, a reservoir filling because of rainfall.

We let $\mathbf{x} = (x_1, \dots, x_S)$ denote a vector of unit states (the system state), $\mathbf{a} = (a_1, \dots, a_S)$ a vector of control decisions (a system action), $A(\mathbf{x}, \psi) = A_1(x_1, \psi) \times \dots \times A_S(x_S, \psi)$ the set of feasible system actions, $p(\mathbf{a}) = \sum_{s=1}^S p_s(a_s)$ the total power produced, $c(\mathbf{x}, \mathbf{a}, \psi) = \sum_{s=1}^S c_s(x_s, a_s, \psi)$ the total cost, and $\chi_t(\mathbf{x}, \mathbf{a}, \psi) = (\chi_{1,t}(x_1, a_1, \psi), \dots, \chi_{S,t}(x_S, a_S, \psi))$ the corresponding vector of next-period unit states.

The DM chooses actions in each period with the goal of maximizing the expected total reward (or minimizing total costs) over the horizon, subject to the constraint of meeting demand in each period and in each state. For ease of later interpretation, we maximize rewards rather than minimize costs. We formulate this problem as a DP: taking the terminal value $V_{T+1}^*(\mathbf{x}, \psi_{T+1}) = 0$,² we can write the optimal value function for earlier periods as

$$\begin{aligned} V_t^*(\mathbf{x}, \psi_t) &= \max_{\mathbf{a} \in A(\mathbf{x}, \psi_t)} -c(\mathbf{x}, \mathbf{a}, \psi_t) \\ &\quad + \mathbb{E}[V_{t+1}^*(\chi_t(\mathbf{x}, \mathbf{a}, \psi_t), \tilde{\psi}_{t+1}) \mid \psi_t] \quad (1) \\ \text{s.t. } p(\mathbf{a}) &= d_t(\psi_t). \end{aligned}$$

Here, $d_t(\psi_t)$ is the demand in period t given world-state ψ_t , and $\mathbb{E}[- \mid \psi_t]$ denotes the expectation over the next-period world-state $\tilde{\psi}_{t+1}$, conditioned on the current world-state ψ_t . As we will discuss in Section 4.2, we assume that the system includes units that shed unmet demand or unneeded supply; these ensure the DP (1) will have a feasible solution. We also assume that the optimal value in any given state will be attained by some vector of actions \mathbf{a} . Note that this formulation assumes that there are no network constraints in the system; we consider the network case in Electronic Companion Section EC-4.

3. A Lagrangian Relaxation

The total costs and power produced in the system model are sums of the costs and production of individual units, but the DP (1) is complicated by the constraint that requires total production to equal demand in every period and every state. These constraints link decisions

and actions across units. In this section, we describe a Lagrangian model that relaxes these linking constraints by introducing stochastic Lagrange multipliers or “prices” associated with these constraints. In this section, we first develop the theory of this Lagrangian model and then discuss some example price models.

3.1. The Lagrangian Model

In the Lagrangian relaxation, we “dualize” the linking constraints in (1) that require production to match demand by introducing Lagrange multipliers $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_T)$ for these constraints. For now, we will allow the period- t Lagrange multiplier λ_t to be any function that is measurable with respect to \mathcal{F}_t ; that is, $\lambda_t: \Psi^t \rightarrow \mathbb{R}^1$. The Lagrange multiplier process $\boldsymbol{\lambda}$ is thus adapted to the filtration \mathbb{F} and $\boldsymbol{\lambda}: \Psi^1 \times \dots \times \Psi^T \rightarrow \mathbb{R}^T$. We let $\eta_t = (\psi_0, \dots, \psi_t)$ denote the history of world-states up to period t , so the period- t Lagrange multipliers can be viewed as a function $\lambda_t(\eta_t)$ of this world-state history.

In our numerical work, we will restrict the Lagrange multipliers $\boldsymbol{\lambda}$ to be in a set Λ with a simpler form, for example, constant functions or linear functions of demand. Although we assume the world-state process is Markovian, the Lagrange multiplier processes need not be. For example, an optimal Lagrange multiplier (or price for power) in a world-state with medium demand may be higher if the preceding demands were low than they would be if the preceding demand states were high, reflecting the need to induce units to start or shut down to meet the current demand.

Taking the terminal value to be $L_{T+1}^\Lambda(\mathbf{x}, \eta_{T+1}) = 0$, the Lagrangian system DP can be written

$$\begin{aligned} L_t^\Lambda(\mathbf{x}, \eta_t) &= \max_{\mathbf{a} \in A(\mathbf{x}, \psi_t)} \lambda_t(\eta_t)(p(\mathbf{a}) - d_t) - c(\mathbf{x}, \mathbf{a}, \psi_t) + \\ &\quad + \mathbb{E}[L_{t+1}^\Lambda(\chi_t(\mathbf{x}, \mathbf{a}, \psi_t), (\eta_t, \tilde{\psi}_{t+1})) \mid \psi_t]. \end{aligned} \quad (2)$$

Here, the dependence of the Lagrangian value functions on the world-state history η_t , rather than the current world-state ψ_t (as in (1)), reflects the most general form of Lagrange multipliers.

This Lagrangian has some nice properties and can be decomposed into unit-specific value functions, as described in the following proposition. These results are standard in the loosely coupled DP literature (see, e.g., Hawkins 2003, Adelman and Mersereau 2008, Brown and Smith 2020), though in this earlier work, the Lagrange multipliers are assumed to be constant in each period rather than arbitrary functions of the history of world-states. Similar results appear in Brown and Zhang (2022).

Proposition 1 (Properties of the Lagrangian Value Function). *For any $\boldsymbol{\lambda}$, $\eta_1 = (\psi_1)$, and \mathbf{x} ,*

(a) *Decomposition:*

$$L_1^\lambda(\mathbf{x}, \eta_1) = -\sum_{t=1}^T \mathbb{E}[\lambda_t(\tilde{\eta}_t) d_t(\tilde{\psi}_t)] + \sum_{s=1}^S V_{s,1}^\lambda(x_s, \eta_1) \quad (3)$$

where the unit-specific value functions $V_{s,1}^\lambda(x_s, \eta_1)$ are given by the DP recursion

$$V_{s,t}^\lambda(x_s, \eta_t) = \max_{a_s \in \mathcal{A}_s(x_s, \psi_t)} \lambda_t(\eta_t) p_{s,t}(a_s) - c_{s,t}(x_s, a_s, \psi_t) + \mathbb{E}[V_{s,t+1}^\lambda(x_{s,t}(x_s, a_s, \psi_t), (\eta_t, \tilde{\psi}_{t+1})) | \psi_t] \quad (4)$$

with terminal case $V_{s,T+1}^\lambda(x_s, \eta_{T+1}) = 0$.

(b) *Upper bound:* $V_1^*(\mathbf{x}, \psi_1) \leq L_1^\lambda(\mathbf{x}, \eta_1)$.

(c) *Convexity:* $V_{s,1}^\lambda(x_s, \eta_1)$ and $L_1^\lambda(\mathbf{x}, \eta_1)$ are piecewise linear and convex in λ .

The decomposition result in Proposition 1(a) follows from rearranging terms in (2). The fact that the Lagrangian DP decomposes in this way means the difficulty of evaluating the system Lagrangian (3) is determined by the complexity of the unit-specific DPs (4). Intuitively, in these unit-specific DPs, the units are paid a price λ_t per unit of power produced in period t , and each unit is managed to maximize its own expected reward. We will consider these unit-specific DPs in more detail in Section 4 and discuss properties that simplify their solution.

Proposition 1(b) follows because the Lagrangian DP (2) is a relaxation of the original DP (1): in (2), the DM could choose actions that are feasible and optimal for (1) and earn the same value as (1) but could also choose actions that do not match demand and possibly earn more. The Lagrangian value function thus provides an upper bound on the primal DP (1). Naturally, we would like to find the best such bound. Given a set of feasible Lagrange multipliers λ and an initial state (\mathbf{x}, ψ) , we consider a (constrained) *Lagrangian dual problem*:

$$\min_{\lambda \in \Lambda} L_1^\lambda(\mathbf{x}, \eta_1). \quad (5)$$

If the set of allowed Lagrange multiplier functions Λ is convex, Proposition 1(c) says this Lagrangian dual problem is a convex optimization problem.

To study the Lagrangian dual problem (5) in more detail, it is helpful to understand the gradient structure of $L_1^\lambda(\mathbf{x}, \eta_1)$. Toward this end, we fix the initial state (\mathbf{x}, η_1) , and let $L(\lambda) = L_1^\lambda(\mathbf{x}, \eta_1)$ and $V_s(\lambda) = V_{s,1}^\lambda(x_s, \eta_1)$. Let α_s denote an \mathbb{F} -adapted policy for managing unit s that selects actions in each period that are feasible (in $\mathcal{A}_s(x_s, \psi_t)$) for the resulting sequence of unit states and the exogenous sequence of world-states; we let \mathcal{A}_s denote the set of all such feasible policies for unit s . Let $c_{s,t}(\alpha_s, \eta_t)$ and $p_{s,t}(\alpha_s, \eta_t)$ denote the period- t cost and energy production at unit s given policy α_s and world-state history η_t . With this notation, we can write the

unit-specific DP (4) as a maximization of the total reward over unit-specific policies α_s :

$$V_s(\lambda, \alpha_s) = \sum_{t=1}^T \mathbb{E}[\lambda_t(\tilde{\eta}_t) p_{s,t}(\alpha_s, \tilde{\eta}_t) - c_{s,t}(\alpha_s, \tilde{\eta}_t)],$$

$$V_s(\lambda) = \max_{\alpha_s \in \mathcal{A}_s} V_s(\lambda, \alpha_s). \quad (6)$$

We can characterize the gradients of the unit-specific DPs and Lagrangian as follows. Suppose we view the Lagrange multiplier process $\lambda = (\lambda_1(\eta_1), \dots, \lambda_T(\eta_T))$ as a vector

$$(\lambda_1(\eta_{1,1}), \dots, \lambda_1(\eta_{1,n_1}), \dots, \lambda_T(\eta_{T,1}), \dots, \lambda_T(\eta_{T,n_T})),$$

where $\eta_{t,1}, \dots, \eta_{t,n_t}$ denotes the possible world-state histories in period t , and $n_t = |\Psi|^t$. The set of feasible Lagrange multiplier processes Λ can then be viewed as a subset of \mathbb{R}^N where $N = \sum_{t=1}^T n_t$. If we view the evolution of world-states as a (nonrecombining) scenario tree, N is the total number of nodes in the tree. Let $\pi(\eta_{t,i})$ denote the probability of world-state history $\eta_{t,i}$ occurring, and let

$$\mathbf{d}_\pi = (\pi(\eta_{1,1})d_1(\eta_{1,1}), \dots, \pi(\eta_{1,n_1})d_1(\eta_{1,n_1}), \dots, \pi(\eta_{T,1})d_T(\eta_{T,1}), \dots, \pi(\eta_{T,n_T})d_T(\eta_{T,n_T}))$$

denote the probability-weighted demand realizations corresponding to these periods and world-state histories.

Proposition 2 (Gradients of the Lagrangian). *Let $\mathcal{A}_s^*(\lambda)$ denote the set of optimal policies for unit s given Lagrange multiplier process λ .*

(a) *Subgradients for the unit-specific problems:* For any $\alpha_s \in \mathcal{A}_s^*(\lambda)$,

$$\nabla_s(\alpha_s) = (\pi(\eta_{1,1})p_{s,1}(\alpha_s, \eta_{1,n_1}), \dots, \pi(\eta_{1,n_1})p_{s,1}(\alpha_s, \eta_{1,n_1}), \dots, \pi(\eta_{T,1})p_{s,T}(\alpha_s, \eta_{T,1}), \dots, \pi(\eta_{T,n_T})p_{s,T}(\alpha_s, \eta_{T,n_T})) \quad (7)$$

is a subgradient of V_s at λ ; that is,

$$V_s(\lambda') \geq V_s(\lambda) + \nabla_s(\alpha_s)^\top (\lambda' - \lambda) \quad \text{for all } \lambda'$$

The subdifferential (the set of all subgradients) of V_s at λ is

$$\partial V_s(\lambda) = \mathbf{conv}\{\nabla_s(\alpha_s) : \alpha_s \in \mathcal{A}_s^*(\lambda)\}$$

where $\mathbf{conv}A$ denotes the convex hull of the set A .

(b) *Subgradients for the Lagrangian:* The subdifferential of L at λ is

$$\partial L(\lambda) = -\mathbf{d}_\pi + \sum_{s=1}^S \partial V_s(\lambda) \quad (8)$$

$$= -\mathbf{d}_\pi + \mathbf{conv}\left\{\sum_{s=1}^S \nabla_s(\alpha_s) : \alpha_s \in \mathcal{A}_s^*(\lambda) \quad \forall s\right\}, \quad (9)$$

where the sums are setwise (i.e., Minkowski) sums.

Proposition 1(a) of this result follows from Danskin's theorem (see, e.g., Bertsekas et al. 2003, proposition

4.5.1); Proposition 1(b) then follows from standard results in convex analysis. Again, these are standard results in the loosely coupled DP literature, adapted here to the general form of \mathbb{F} -adapted Lagrange multipliers. For completeness, we provide a proof in Section EC-1 in the electronic companion.

The following theorem provides optimality conditions for the Lagrangian dual problem (5).

Theorem 1 (Optimality Conditions). *Suppose Λ is a convex set of feasible Lagrange multiplier processes.*

(a) *Abstract form: $\lambda^* \in \Lambda$ is an optimal solution for the Lagrangian dual problem (5) if and only if*

$$\mathbf{0} \in \partial L(\lambda^*) + \mathcal{N}_\Lambda(\lambda^*), \quad (10)$$

where $\mathcal{N}_\Lambda(\lambda^*)$ is the normal cone of Λ at λ^* . Given the (assumed) convexity of Λ , $\mathcal{N}_\Lambda(\lambda^*)$ is the set of all $z \in \mathbb{R}^N$ such that $z^\top(\lambda - \lambda^*) \leq 0$ for all $\lambda \in \Lambda$.

(b) *Mixture form: Let $\mathcal{A}_s^*(\lambda)$ denote the set of optimal policies for unit s given Lagrange multiplier process λ . Then, λ^* is an optimal solution for the Lagrangian dual problem (5) if and only if, for each s , there is a set of policies $\{\alpha_{s,i}\}_{i=1}^{m_s}$ with $\alpha_{s,i} \in \mathcal{A}_s^*(\lambda^*)$ and mixing weights $\{\gamma_{s,i}\}_{i=1}^{m_s}$ ($\gamma_{s,i} > 0$ and $\sum_{i=1}^{m_s} \gamma_{s,i} = 1$) such that*

$$d_\pi - \sum_{s=1}^S \sum_{i=1}^{m_s} \gamma_{s,i} \nabla_s(\alpha_{s,i}) \in \mathcal{N}_\Lambda(\lambda^*), \quad (11)$$

where $\nabla_s(\alpha_{s,i})$ is the probability-weighted time- and state-contingent production vector for policy $\alpha_{s,i}$, for unit s , given as the gradient (7).

(c) *Linear price functions: Suppose Λ is contained in a K -dimensional linear subspace of \mathbb{R}^N with*

$$\lambda_t(\eta_t) = \beta_1 b_{1,t}(\eta_t) + \cdots + \beta_K b_{K,t}(\eta_t), \quad (12)$$

that is, a linear combination of some set of “basis functions” $b_{k,t}(\eta_t)$. Then, λ^* is an optimal solution for the Lagrangian dual problem if and only if, for each s , there is a set of optimal set of policies $\{\alpha_{s,i}\}_{i=1}^{m_s}$ with $\alpha_{s,i} \in \mathcal{A}_s^*(\lambda^*)$ and mixing weights $\{\gamma_{s,i}\}_{i=1}^{m_s}$ such that, for all k ,

$$\sum_{s=1}^S \sum_{i=1}^{m_s} \gamma_{s,i} \sum_{t=1}^T \mathbb{E}[b_{k,t}(\tilde{\eta}_t) p_{s,t}(\alpha_{s,i}, \tilde{\eta}_t)] = \sum_{t=1}^T \mathbb{E}[b_{k,t}(\tilde{\eta}_t) d_t(\tilde{\eta}_t)] \quad (13)$$

with $\sum_{s=1}^S m_s \leq S + K$; thus, no more than K units will have mixed policies.

We provide a proof in Section EC-1 in the electronic companion. The first result here is a standard result for convex optimization (see, e.g., Bertsekas et al. 2003, proposition 4.7.2). The mixture form in Theorem 1(b) then follows from (9) and (10) using Caratheodory’s representation of the convex set in (9). Noting that $\nabla_s(\alpha_{s,i})$ is the probability-weighted production vector given policy $\alpha_{s,i}$ (optimal for λ^*), the left side of (11) can be interpreted as the (probability-weighted) residual of optimally “fitting” the Lagrangian production model to

the observed demands d_π : the optimality condition (11) requires this residual to lie in the normal cone $\mathcal{N}_\Lambda(\lambda^*)$ of the constraint set Λ . The refinement of Theorem 1(c) follows from Theorem 1(b) and standard results about the form of basic feasible solutions for linear systems of equations. Note that though optimal solutions for the Lagrangian relaxations may involve mixed policies, the feasible dispatch policies we consider (see Section 5) do not involve mixing.

With price models of the form (12), we can work in a K -dimensional subspace of Λ , and the results above imply the Lagrangian dual problem has a particularly nice structure.

Proposition 3 (Lagrangian Dual with Linear Prices). *Suppose the price model $\lambda(\beta)$ has the form (12) with weight vector $\beta \in \mathbb{R}^K$. The Lagrangian dual problem (5) can be written*

$$\min_{\beta} L_1^{\lambda(\beta)}(x, \eta_1), \quad (14)$$

where

(a) *the objective $L_1^{\lambda(\beta)}$ is piecewise linear and convex in β , and*

(b) *for any policy α_s that is optimal for the unit-specific DP (4) with prices $\lambda(\beta)$, $\nabla(\alpha) = (\nabla_1, \dots, \nabla_K)$, where*

$$\nabla_k = - \sum_{t=1}^T \mathbb{E}[b_{k,t}(\tilde{\eta}_t) d_t(\tilde{\eta}_t)] + \sum_{s=1}^S \sum_{i=1}^T \mathbb{E}[b_{k,t}(\tilde{\eta}_t) p_{s,t}(\alpha_s^{\lambda(\beta)}, \tilde{\eta}_t)] \quad (15)$$

is a subgradient of $L_1^{\lambda(\beta)}$ at β .

Given this piecewise linear convex structure, it is natural to use cutting-plane methods to solve (14), as in this setting, the cutting-plane algorithm is guaranteed to converge to an optimal λ in a finite number of iterations. The objective and gradient calculations both decompose across units, and the unit-specific gradient terms (inside the sum over s in (15)) are straightforward to calculate when solving the unit-specific DPs (4). This implies that the values and gradients may be evaluated in parallel within the cutting-plane routine, which may be convenient for large systems (i.e., large S). See, for example, Brown and Smith (2020) for more discussion of the use of cutting-plane techniques for these kinds of problems.

We can better understand the results of Theorem 1 by considering specific examples of constraint sets Λ for the Lagrange multiplier or price process.

3.2. No Constraints on the Price Model

First, suppose we have no constraints on the form of the price process other than that it is \mathbb{F} -adapted. In this case, $\Lambda = \mathbb{R}^N$, where, as noted above, $N = \sum_{t=1}^T n_t$ is the total number of world-state histories over time. In this case, $\mathcal{N}_\Lambda(\lambda^*)$ includes only the zero vector, and (11) means that the optimal Lagrange multiplier process λ^* leads to

a mixture of policies that exactly matches the realized demand in every possible world-state in every period. These mixed policies “convexify” the underlying objective function, and thus, $L(\lambda^*)$ need not provide a tight bound on the optimal value with a feasible policy; that is, there may be a duality gap.

The following theorem shows that, with no constraints on the price model, this duality gap is independent of the number of units in the system (S) and the number of scenarios considered (N). The proof of this result involves characterizing the “lack of convexity” in the model that is exploited by mixing policies in the Lagrangian. Toward this end, we will say a unit’s production possibilities are convex if, for each $a'_s, a''_s \in A_s$ and $\mu \in [0, 1]$, there exists an $a_s^\mu \in A_s$ such that $p(a_s^\mu) = \mu p(a'_s) + (1 - \mu)p(a''_s)$.³

Theorem 2 (Unconstrained Duality Gap). *Suppose the production possibilities are convex for all units, and unit costs and production are bounded. Then, for any x and $\eta_1 (= \psi_1)$ and optimal Lagrange multiplier process λ^* for the Lagrangian dual problem (5) with $\Lambda = \mathbb{R}^N$, we have*

$$L_1^*(x, \eta_1) - V_1^*(x, \psi_1) \leq \kappa,$$

where, for a fixed time horizon, κ is independent of both N and S .

Assuming optimal costs grow linearly in the number S of units, Theorem 2 implies that the relative duality gap for the unconstrained price model converges to zero at rate $1/S$ as the number of units S grows large. Thus, we would expect these duality gaps to be relatively small for large systems. The key insight underlying this result is that with a single demand-matching constraint for each period and state, the optimal Lagrangian policy will mix policies for a single unit in each period/state. Of course, different units may mix in different periods and states, and the bound is given by considering the largest potential gap across the S units. The constant gap in Theorem 2 stands in contrast to $O(N)$ the duality gap of Takriti and Birge (2000) for scenario decomposition and the $O(\sqrt{S})$ gaps of Brown and Smith (2020) and Brown and Zhang (2022) for Lagrangian relaxations in different settings. We provide a proof of Theorem 2 in Section EC-1 in the electronic companion.

Though solving the unconstrained Lagrangian dual may be difficult when the number of world-state histories is large, Theorem 2 suggests that in large systems with complex uncertainties, little is lost in working with the Lagrangian relaxation, and with judiciously chosen price models, the LR relaxation (2) can provide a good approximation of the original value function (1).

3.3. Period-Constant Prices

Next, suppose that Λ consists of Lagrange multipliers that are constant in each period, as is common in the

weakly coupled DP literature, for example, Hawkins (2003), Adelman and Mersereau (2008), and Brown and Smith (2020). This case can be represented in the form of (12) by considering T basis functions $b_{k,t}(\eta_t) = 1$ for $k = t$ and all η_t , and zero otherwise. The optimality condition (13) then becomes

$$\sum_{s=1}^S \sum_{i=1}^{m_s} \gamma_{s,i} \mathbb{E}[p_{s,i}(\alpha_s, \tilde{\eta}_t)] = \mathbb{E}[d_t(\tilde{\eta}_t)] \text{ for all } t, \quad (16)$$

where the $\gamma_{s,i}$ are mixing weights in Theorem 1(c). These conditions can be interpreted as saying that the optimal Lagrange multipliers (prices) are such that total production in the Lagrangian matches demand “in expectation” in each period, where the expectation includes the mixing of policies as well as uncertainty about demand. Note that these prices are deterministic and reflect time-of-day effects but are not sensitive to changing demand or world-states. In Section 6, we will see that this deterministic price model does not perform very well in our numerical experiments.

3.4. Period-Linear Prices

Now suppose that Λ consists of Lagrange multipliers in each period that are affine functions of demand in that period: $\lambda_t(\eta_t) = \beta_{t,0} + \beta_{t,1}d_t(\eta_t)$. This model can be represented in the form of (12) using $2T$ basis functions. The first T basis functions are constants $b_{k,t}(\eta_t) = 1$ for all η_t , and zero otherwise, as above; the corresponding β_k represent period-specific intercept terms. The next T basis functions are the period- t demands, $b_{k,t}(\eta_t) = d_t(\eta_t)$ for $k = T + t$, and zero otherwise; the corresponding β_k represent period-specific slope terms. The optimality conditions (13) then require that production in the Lagrangian matches demand in expectation in each period, as in (16) above, and

$$\sum_{s=1}^S \sum_{i=1}^{m_s} \gamma_{s,i} \mathbb{E}[p_{s,i}(\alpha_s, \tilde{\eta}_t)d_t(\tilde{\eta}_t)] = \mathbb{E}[d_t(\tilde{\eta}_t)^2] \text{ for all } t. \quad (17)$$

Conditions (16) and (17) imply that with optimal weights (b_t in (12)), the simple linear regression of total production in period t in the Lagrangian model regressed on demand in period t would have unit slope and zero intercept in each period. This price model is stochastic, captures time-of-day effects, and has time of day-specific demand effects. In Section 6, we will see that this price model performs very well in our numerical experiments.

3.5. Other Price Models

Of course, many other forms of price models are possible. One could always add other basis functions to the Lagrange multiplier model. For example, we could consider a model that includes lagged demand as well as the current demand in the Lagrangian model to reflect

the need to induce units that were previously off to start or units that were on to shut down to meet the current demand (as discussed in Section 3.1). However, we want to choose a functional form for the price model with computational considerations in mind so the Lagrangian dual problem (5) can be solved within a reasonable time. Specifically, we want to have (i) relatively few parameters to be optimized, (ii) gradients that are “easy” to calculate, and (iii) small world-states and limited dependence on the history of world-states, so the state spaces in the unit-specific DPs are not too large. With many periods and significant uncertainty, the number of parameters (N) in the unconstrained case will likely be prohibitively large.

In practice, we recommend proceeding incrementally, starting with a simple model and then complicating as necessary. In our numerical experiments, we have found good results with price models that are quite simple. For example, there is no a priori reason that we need to have period-specific constant or linear terms in the price model, as assumed in Sections 3.3 and 3.4. One could, for example, assume prices in any period are a linear function of demand in that period, with the same slope and intercept in every period; this would correspond to a price model with just two parameters. In Section 6, we report results for a simple price model with seven parameters; we discuss the choice of these basis functions in Section 6.2. Choosing the form of the price model (i.e., choosing basis functions) is something of an art and, at a high level, similar to choosing basis functions for a regression model or other approximate DP methods.

4. Unit Models and Properties

We now describe the units we will consider as part of our example power system. For each unit, we describe its state space X_s , the possible actions $A_s(x_s, \psi_t)$, costs $c_s(x_s, a_s, \psi_t)$, and production $p_s(a_s)$ levels, where these functions may depend on the current unit state x_s , action a_s , and world-state ψ_t . We focus on the kinds of units that are included in our example system. In that data set, there are 127 generation units (Section 4.1) and two pumped-storage hydro units (Section 4.3). We discuss demand- and supply-shedding units in Section 4.2.

4.1. Generation Units

The broad category of generation units includes nuclear plants, thermal generation plants, peaking plants, etc., with parameters and constraints varying by unit. For example, nuclear plants have large startup and shutdown costs and low variable costs. Peaking plants (e.g., natural gas turbines) have low startup and shutdown costs but typically have relatively high operating costs. Thermal plants are between these two cases, having significant startup costs and ramping constraints but potentially having lower operating costs. We consider

these various kinds of units as parametric variations within the class of generation units.

A generation unit may be “off”—a state we denote by the number zero—or “on” and producing at some production level p between lower bound $p_s > 0$ and upper bound \bar{p}_s ; thus, the state space X_s is $\{0\} \cup [p_s, \bar{p}_s]$. The action space A_s represents the next-period state (off or producing); thus, $A_s = X_s$ and $\chi_{s,t}(x_s, a_s, \psi_t) = a_s$. Production may ramp up or down in a given period by at most $r_{s,u}$ and $r_{s,d}$ (respectively); the maximum production achievable at startup is $p_{s,u}$, and the maximum production rate where shutdown is possible is $p_{s,d}$. Thus, the set of feasible actions for a generation unit s is

$$A_s(x_s, \psi_t) = \begin{cases} \{0\} \cup [p_s, p_{s,u}] & \text{if } x_s = 0, \\ \{0\} \cup ([p_s, \bar{p}_s] \cap [x_s - r_{s,d}, x_s + r_{s,u}]) & \text{if } x_s \neq 0 \text{ and } x_s \leq p_{s,d}, \\ [p_s, \bar{p}_s] \cap [x_s - r_{s,d}, x_s + r_{s,u}] & \text{otherwise.} \end{cases}$$

The production for generation unit s is $p_s(a_s) = a_s$.⁴

The production costs include fixed costs $c_{s,n}$ (sometimes called “no-load costs”), as well as startup costs $c_{s,u}$ and shutdown costs $c_{s,d}$. We can write the cost function as

$$c_s(x_s, a_s, \psi_t) = \begin{cases} 0 & \text{if } x_s = 0 \text{ and } a_s = 0, \\ c_{s,u} + c_{s,n} + c_{s,f}a_s & \text{if } x_s = 0 \text{ and } a_s \neq 0, \\ c_{s,d} & \text{if } x_s \neq 0 \text{ and } a_s = 0, \\ c_{s,n} + c_{s,f}a_s & \text{otherwise.} \end{cases}$$

In principle, we could have world-state dependent variable costs $c_{s,f}$ (or other costs) representing uncertainty in fuel costs. However, in our numerical examples, we will assume constant fuel costs.

Though the state and action spaces here include a continuum of possible production levels, in the context of the unit-specific DPs for the Lagrangian, we can restrict our focus to a finite set of production levels that are potentially optimal. The intuition behind this result generalizes the idea of a “bang-bang” solution: if we had no ramping constraints or startup or shutdown limits or costs, it would be optimal for a plant to be producing at either its upper limit or at zero, according to whether the operating profit, $\bar{p}_s(\lambda_t(\eta_t) - c_{s,f}) - c_{s,n}$, is positive or negative for a given world-state. With ramping constraints or startup and shutdown constraints and costs, the DM needs to consider the effects of current actions on future feasible sets of actions.

This logic suggests that there is a discrete set of production levels that an optimal policy would visit. These “grid” values are

$$\mathcal{G}_s = \{\hat{x} + ir_{s,u} - jr_{s,d} : i, j \in \mathbb{Z}_+\} \cap [p_s, \bar{p}_s]. \quad (18)$$

In words, these are the feasible values (between \underline{p}_s and \bar{p}_s) that can be reached by ramping up or down as much as possible ($r_{s,u}$ and $r_{s,d}$) repeatedly over time, starting from the unit's initial production level ($p_{s,1}$ if not off initially), the minimum or maximum levels \underline{p}_s and \bar{p}_s , the maximum startup level $p_{s,u}$, or the maximum shutdown level $p_{s,d}$. For instance, one of the generation units in our numerical examples has $\underline{p}_s = 310$, $\bar{p}_s = 844$, $p_{s,u} = p_{s,d} = 310$, and ramp limits $r_{s,u} = r_{s,d} = 180$. If the unit starts in the off state, the grid values are $\{310, 490, 670, 844, 664, 484\}$. The following result says that rather than considering a continuum of possible states in the unit-specific DP, we can focus on a state space consisting of these six production levels plus the off-state.

Proposition 4 (Properties of Generation Units). *Consider any $\lambda \in \Lambda$, and let \mathcal{G}_s denote the set of grid points defined in (18) for generation unit s . For every t and η_t , the following hold:*

(a) *Without shutdown decisions (i.e., if $c_{s,d} = \infty$), $V_{s,t}^\lambda(x, \eta_t)$ is piecewise linear and concave on $x \in [\underline{p}_s, \bar{p}_s]$ with breakpoints in \mathcal{G}_s .*

(b) *With shutdown decisions, $V_{s,t}^\lambda(x, \eta_t)$ is convex in x between grid points: that is, suppose g_1 and g_2 are adjacent points in \mathcal{G}_s , $\mu \in [0, 1]$, and x_1, x_2 , and x_μ are such that $x_\mu = \mu x_1 + (1 - \mu)x_2$ and $g_1 \leq x_1 \leq x_\mu \leq x_2 \leq g_2$. Then, $V_{s,t}^\lambda(x_\mu, \eta_t) \leq \mu V_{s,t}^\lambda(x_1, \eta_t) + (1 - \mu)V_{s,t}^\lambda(x_2, \eta_t)$.*

In both cases, if x is in \mathcal{G}_s and it is optimal to produce, then there exist optimal production levels in \mathcal{G}_s in period t and thereafter.

The result of Proposition 4(a) generalizes the bang-bang logic to incorporate ramping limits and can be proven by backward induction: a DM facing a piecewise linear convex continuation value with rewards that are linear in the production level will either want to ramp up or down as much as possible (which, if starting at a grid point, would be a grid point) or else operate at a breakpoint of the continuation value (which, by the induction hypothesis, is also a grid point). Though it is not surprising that including shutdown decisions would destroy the concavity of value function, it is perhaps surprising that local *convexity* emerges. This local convexity is enough to ensure the desired result: with rewards that are linear in the production level, producing at any point between grid points is dominated by producing at one of the bracketing grid points. In practice, these sets of grid points may be quite small. For instance, with the 129 generation units in the Duke Energy system modeled in Section 6, the number of grid points ranges from 1 (when $\underline{p}_s = \bar{p}_s$) to 10, with an average of 2.60.⁵

Figure 2(a) illustrates Proposition 4(b), showing a value function for a generation unit with shutdown decisions. For this unit, $\underline{p}_s = 0.1$, $\bar{p}_s = 1.0$, and $r_{s,u} = r_{s,d} = 0.1$, $p_{s,u} = 0.2$, and $p_{s,d} = 0.3$; the grid points (shown as dots in the figure) thus lie at increments of

0.1. The convexity of the value function between grid points is evident in several places.⁶

4.2. Demand- and Supply-Shedding Units

As mentioned in Section 2, we assume that there are units for shedding excess demand and supply. For example, the demand-shedding unit may reflect a “rolling blackout” in the case that demand exceeds supply. These are special cases of the production units described in the previous subsection. Intuitively, one can think of the demand-shedding unit as a peaking plant with no fixed costs, startup costs, or ramping limits and marginal cost ($c_{s,f}$) equal to the “value of lost load.” The supply-shedding unit is a peaking unit that produces negative power; the fuel costs $c_{s,f}$ for the supply-shedding unit are the *negative* of the cost of shedding unneeded supply. Both units are unconstrained in terms of their ability to shed: (i) for the demand-shedding unit, $\underline{p}_s = 0$, and \bar{p}_s is a large positive value that serves as an upper bound on system demand, and (ii) for the supply-shedding unit, \underline{p}_s is a large negative value, and $\bar{p}_s = 0$. In the Lagrangian relaxation, it is optimal for these units to produce \bar{p}_s when $\lambda_t(\eta_t) \geq c_{s,f}$ and optimal to produce \underline{p}_s when $\lambda_t(\eta_t) \leq c_{s,f}$. Finite amounts of shedding are achieved in an optimal solution to the Lagrangian by taking $\lambda_t(\eta_t) = c_{s,f}$ and mixing (as in Theorem 1) between the maximum and minimum production levels, which are both optimal at this price.


In addition to these “last ditch” shedding units, one could consider a number of other kinds of shedding units with different costs and allowed quantities. The general framework is quite flexible in this regard. Note that in our numerical experiments, our recommended feasible dispatch policies (described in Section 5) do not require any demand or supply-shedding in any scenario, though other policies we consider do.

4.3. Storage Units

The state x_s for a storage unit represents the energy storage level. Storage must be between minimum and maximum levels \underline{x}_s and \bar{x}_s , and thus, $X_s = [\underline{x}_s, \bar{x}_s]$. Actions a_s correspond to the amount of power generated (if positive) or consumed (if negative), so $p(a_s) = a_s$, where the action space $A_s = [\underline{p}_s, \bar{p}_s]$ corresponds to the storage unit limits on power consumption and generation. Here, power consumption corresponds to increasing the energy storage level for potential later generation. The state transition function is

$$\chi_{s,t}(x_s, a_s, \psi_t) = x_s - \gamma_s^- a_s^- - \gamma_s^+ a_s^+,$$

where a_s^+ and a_s^- correspond to the positive and negative parts of a_s , respectively. The constants γ_s^- and γ_s^+ describe the efficiency of the storage unit: every unit of power consumed by the unit increases the energy stored by γ_s^- units, and every unit of power generated requires γ_s^+ units of stored energy. Thus, the action sets for

Figure 2. (Color online) Illustrative Unit-Specific Value Functions

storage unit s are

$$A_s(x_s, \psi_t) = [p_s, \bar{p}_s] \cap [(x_s - \bar{x}_s)/\gamma^-, (x_s - \underline{x}_s)/\gamma^+].$$

There are no costs associated with storage, and in the unit-specific DP, the rewards are $\lambda_t(\eta_t)a_s$, reflecting the value of the power produced or consumed.

As with the generation units without shutdown decisions, the value functions for the storage units can be shown to be piecewise linear concave with breakpoints and optimal production levels restricted to a finite grid that may be determined in advance. Like the generation unit case, these grid points are defined by ramping storage up and down by integer multiples of $\gamma_s^- \bar{p}_s^-$ and $\gamma_s^+ \bar{p}_s^+$, respectively, starting from the initial storage level x_1 and the maximum and minimum storage levels \underline{x}_s and \bar{x}_s . In the Duke energy system modeled in Section 6, there are 83 and 95 grid points for the two storage units. Figure 2(b) shows an example of a value function for a storage unit, with the dots representing possible grid points.

This storage model could be extended to include (random or deterministic) inflows or outflows representing, for example, rainfall or evaporation at a pumped-storage hydro unit or leakage at a battery unit. The piecewise linear convex structure of the model would remain, but one may have to consider a more refined grid to incorporate possible inflows or outflows.

5. Forward-Looking Economic Dispatch

The decomposed policies from solving the unit-specific DPs (4) in the Lagrangian decomposition need to be coordinated to find a production plan that meets the realized actual demand in a given period. We envision these dispatch problems being solved hourly (or more frequently), like the ED problem in current practice. Here, however, we use the unit-specific value functions

to incorporate the longer-term effects of these dispatch decisions in the ED model. Specifically, we aim to choose actions (production levels) to solve a version of the original DP (1) where the next-period value function V_{t+1}^* is replaced by the Lagrangian L_{t+1}^λ :

$$\begin{aligned} \max_{a \in A(x)} \quad & -c(x, a, \psi_t) + \mathbb{E}[L_{t+1}^\lambda(\mathbf{X}_t(x, a, \psi_t), (\eta_t, \tilde{\psi}_{t+1})) \mid \psi_t] \\ \text{s.t.} \quad & p(a) = d_t(\psi_t). \end{aligned} \quad (19)$$

We can use any price process λ in $L_t^\lambda(x)$ in this approach but will focus on optimal price processes (i.e., solve the Lagrangian dual problem (5)) for a given form of price model.

The optimization problem (19) can be formulated as a mixed-integer linear program in various ways. Here, we consider a “convex hull” formulation that is convenient for use with the results provided by the unit-specific DPs. Specifically, we introduce decision variables $w_{s,i} \in [0, 1]$ that are weights associated with the feasible actions $a_{s,i} \in A_s(x_s)$ considered in these unit-specific DPs. These actions include the off state (for generation units only) and the actions $a_{s,i}$ corresponding to production levels in \mathcal{G}_s (the set of grid points defined in Section 4) that are reachable from state x_s . For initial states x_s that are not in \mathcal{G}_s , we augment this set to include actions $x_s + r_{s,u}$ (the maximum ramp up) and $x_s - r_{s,d}$ (the maximum ramp down), provided these points are between the minimum and maximum production levels. We similarly augment the sets of possible next-period states for storage units if the initial state is not in \mathcal{G}_s . We let n_s denote the number of such feasible actions for unit s .

The MILP is then

$$\begin{aligned}
 \max_{w_{s,i}} \quad & \sum_{s=1}^S \sum_{i=1}^{n_s} w_{s,i} (-c_s(x_s, a_{s,i}, \psi_t)) \\
 & + \mathbb{E}[V_{s,t+1}^\Lambda(x_s, a_{s,i}, \psi_t, (\eta_t, \tilde{\psi}_{t+1})) | \psi_t] \\
 \text{s.t.} \quad & \sum_{s=1}^S \sum_{i=1}^{n_s} w_{s,i} p_s(a_{s,i}) = d_t(\psi_t) \\
 & \sum_{i=1}^{n_s} w_{s,i} = 1 \quad \forall s \\
 & w_{s,i} \geq 0 \quad \forall s, i \\
 & w_{s,i} \in \{0, 1\} \quad \forall s, i \text{ corresponding to "off" states.}
 \end{aligned} \tag{20}$$

The expected continuation values appearing in the objective here are calculated when solving the Lagrangian dual for all actions corresponding to production levels that are on the grid (in \mathcal{G}_s), but not for those actions that are off-grid (not in \mathcal{G}_s). We estimate the off-grid expected continuation values using linear interpolation. This linear interpolation is exact for storage units because their value functions are piecewise linear convex with breakpoints at these grid points (as discussed in Section 4.3). However, this linear interpolation is an approximation for generator units because, as discussed in Section 4.1 and illustrated in Figure 2(a), the continuation values for generation units may be strictly convex between grid points; thus, the linear interpolation of these value functions may overestimate of the actual unit-specific continuation values. Note that unlike the multiperiod UC problem (see Section EC-2.2 in the electronic companion), this MILP is small: there is at most one binary variable for each generation unit, and the number of continuous variables is less than the sum of the grid sizes for the units.

The solution to this ED problem calls for unit s to produce at the level given by $\sum_{i=1}^{n_s} w_{s,i}^* p_s(a_{s,i})$ for the optimal weights $w_{s,i}^*$ given by (20). Thus, units may operate at levels between the production levels associated with the identified feasible actions $a_{s,i}$ in (20). The binary constraints associated with the off states in (20) ensure that the solution cannot take convex combinations of the off-state (with zero production) and the other production levels. A simple constraint counting argument implies that for any given setting of the binary variables, there are $S + 1$ equality constraints and, hence, at most $S + 1$ positive weights in any basic feasible solution. Thus, at most, one unit will be taking a convex combination of the production levels considered in (20); this unit can be considered the “marginal unit” for the given scenario, and this convex combination ensures supply equals demand. For example, if the system must shed demand in a given scenario, the optimal solution would have positive weights on zero demand shedding and on the assumed upper

bound on shedding (see Section 4.2), and the solution would shed the convex combination of these two levels. More typically, the marginal unit would be a generation unit or storage unit operating at a convex combination of values between adjacent grid points.⁷

A myopic version of the ED problem would take the expected continuation values in (20) to be zero. We can incorporate commitment constraints by requiring the binary variables for the off states for slow-start units to be either zero or one, according to the solution from the UC problem, and similarly constraining the weights for the storage units. See Section EC-2.2 in the electronic companion for further discussion and an alternative formulation of the ED problem.

Our discussion of the ED problem assumes that the timescale matches that of the DP and its Lagrangian relaxation; in our examples, we will assume one-hour time steps. In practice, some system operators solve ED problems on a much more frequent timescale, for example, every five or 15 minutes. Because these problems can be solved very quickly, these forward-looking EDs can be solved with these more refined timescales. Here, if the Lagrangian DP is solved on an hourly timescale, we need to somehow interpolate the expected continuation values appearing in (18) to provide expected continuation values on the more refined scale. For example, we could assume that the end-of-hour expected continuation values apply over the previous hour. Numerical experiments with a five-minute timescale suggest that our dynamic approach continues to perform well.

6. Numerical Experiments

We evaluate the proposed dynamic approach by applying it in examples involving the Duke Energy Carolinas and Duke Energy Progress systems. We focus on eight specific days and consider three different scenarios representing (i) the “current” (2019) system; (ii) a “future” version of this system with four times the current solar production and twice the storage capacity, roughly representing Duke Energy’s future plans for the Carolinas system (Duke Energy 2023); and (iii) this future system with a carbon tax of \$185 per metric ton of carbon dioxide emissions (Rennert et al. 2022). For each day and system scenario, we consider the performance of the proposed approach and compare it to several benchmarks. In this section, we describe the Duke Energy System (Section 6.1), policies considered in our experiments (Section 6.2), details of the simulations (Section 6.3), and results (Sections 6.4 and 6.5).

6.1. The Duke Energy System

The modeled system has 127 generation units with a total capacity of about 40 gigawatts. Specifically, there are 11 nuclear generation units with a total capacity of 11.1 gigawatts (these will always be on), 36 “slow-start”

units—coal or natural gas thermal units—with a total capacity of 17.0 gigawatts, and 80 “fast-start” units—for example, natural gas turbines—with a total capacity of 9.6 gigawatts. The no-load, start-up, shutdown, and variable costs for these generation units vary depending on the technical details of the unit (e.g., heat rates) and the fuel costs. In our experiments, we take fuel costs to be \$3.80 metric million British thermal units (MMBtu) for coal, \$4.00/MMBtu for natural gas, and \$13.50/MMBtu for oil, without carbon taxes. In the system scenario with the carbon tax, these fuel costs are adjusted to reflect the associated carbon dioxide emissions.

Figure 3 shows supply curves for these generation units, with and without carbon taxes. In Figure 3, (a) and (b), there are two “dashes” connected by a faint line for each of the 127 generation units. We use blue-to-green dashes for slow-start units and pink-to-yellow dashes for fast-start units. The lower dash value for each unit is the average cost of production when the unit is running at full power for one hour; using the notation from Section 4.1, this is $c_{s,n}/\bar{p}_s + c_{s,f}$. The widths of these lower dashes are the maximum generation capacity of the unit (\bar{p}_s). We form the “supply curve” by sorting units in increasing order of their full-power average costs and accumulating their capacities. The vertical lines from each of these lower dashes connect to an upper dash with value equal to the average cost of energy when switching the unit from off to its maximum start-up power, $(c_{s,n} + c_{s,u})/\bar{p}_{s,u} + c_{s,f}$; this represents the average cost of power supplied for an hour when the unit first starts up. The widths of these upper dashes are the unit’s maximum start-up power ($\bar{p}_{s,u}$). The fuel source for each unit is shown at the bottom of the

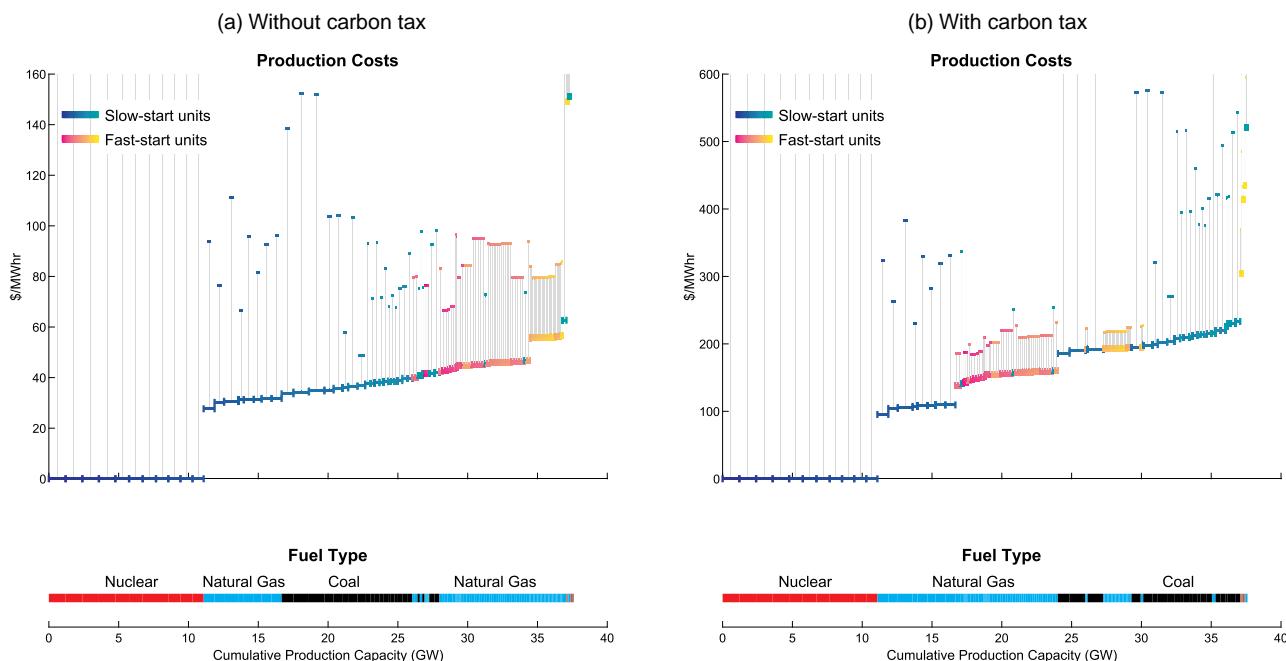
figures. Here, we see that the lowest-cost units by the full-power average cost metric are the nuclear units, which provide energy at no cost; the next lowest-cost units are the natural gas slow-start units. We also see that the slow-start units have larger generation capacities than the fast-start units.

Comparing Figure 3, (a) and (b), we see that, with the carbon tax, the overall costs are higher (note the figures have different y-axes), and the slow-start coal units shift to the right of the natural gas fast-start units in the supply curve. This shift reflects the fact that coal emits about twice as much carbon dioxide as natural gas per unit of energy produced. The carbon tax makes these coal-fired units very expensive to operate; the net effect in our simulations is essentially equivalent to removing the coal units from the system.

These supply curves illustrate the challenge of managing operations under uncertainty. Measured by full-power costs, the lowest-cost units in the 12- to 25-gigawatt range—where much of the variation in system (net) demand lies—are slow-start units: these slow-start units have lower costs when operating at full power but typically have higher start-up costs than the fast-start units. Thus, uncertainty in demands in this range creates a difficult trade-off: do we meet changing demand with slow-start or fast-start units?

In addition to the generation units, the modeled system includes 2.2 gigawatts of production capacity from two pumped-storage hydro units. The Duke Energy system includes both utility-scale solar and distributed (“behind the meter”) solar, which are both netted out from our forecasts of net demand (see Section 6.3). Finally, as discussed in Section 2, the demand-shedding

Figure 3. (Color online) Supply Curves with and Without Carbon Tax



unit can shed demand at a cost of \$3,000/megawatt-hour (the “value of lost load”), and the supply-shedding unit can shed excess supply at \$10/megawatt-hour.⁸

6.2. Policies Considered

In our examples, we simulate seven policies: the proposed dynamic approach with three different price models and four benchmark policies for comparison. The three price models are:

1. Period-constant prices (Section 3.3). This model has 24 parameters, one for each hour.
2. Period-linear prices (Section 3.4). This model has 48 parameters, that is, slope and intercept terms for each hour.
3. A “simple” price model. This model has seven basis functions: a constant (1), the period- t demand d_t , period- t demand minus forecast period- t demand, dummy indicator variables for the first two periods, and two “check” functions, $(d_t - 16)^+$ and $(d_t - 20)^+$, where d_t is measured in gigawatts.

Simulating the system with these dynamic policies requires solving the Lagrangian dual problem (5) for a given price model to determine the optimal parameters of the price model. As discussed after Proposition 3, we use a cutting-plane algorithm to solve these optimization problems. These optimization problems must be solved once for each simulation. Within each simulation, in each sample demand scenario, we use the resulting unit-specific value functions to dynamically dispatch production, as discussed in Section 5.

The first two price models above are natural candidates to consider; the results for the period-constant price model will illustrate the importance of using a stochastic price model. We used a combination of intuition and experimentation to choose the basis functions for the simple price model. First, it seems natural to include constant and demand terms to capture a linear price-to-demand relationship. The demand minus forecast term was added to capture the idea that higher (lower) than expected demands might lead to higher (or lower) prices. The two check functions were added to capture potential nonlinearity (e.g., convexity) in the effective supply curve, noting that higher demands may lead to starting more expensive units (e.g., peakers). The use of these check functions was motivated by the theoretical result that any convex function can be represented by a (positive) linear combination of check functions; the particular thresholds of 16 and 20 gigawatts were selected by inspecting the supply curve (Figure 3(a)) and some experimentation. The dummy variables for the first two hours were also added as the result of some experimentation. We used this same set of basis functions in all of our simulations. Tuning the choice of basis functions for particular days, seasons, or system scenarios could lead to somewhat better results, though, as we will see, there is not much room for improvement in our examples.

We compare the performance of the dynamic policies above with:

4. Unit commitment with myopic dispatch (current practice). Here, we first solve the UC problem for a week with demands set to their expected values. Then, we “commit” the slow-start units to be on or off (by setting the corresponding binary variables) and the storage flows to their UC solutions and solve a myopic ED problem to determine the production levels in each period. (See Section EC-2.1 in the electronic companion for details.)

5. Myopic dispatch without commitment. We solve the myopic ED problem in each period without committing slow-start units or storage.

The first of these policies represents “current practice” at many integrated utilities, as discussed in Section 1. We acknowledge that this policy is a simplification of what system operators actually do: as discussed in Section 1, if the actual demands deviate significantly from the forecasts, the operator may, for example, adjust the unit commitments and storage plans informally “on the fly” as the day unfolds. The operator may also rerun the UC problem with revised forecasts to develop a new plan if it becomes evident that the actual demand is significantly different from the original forecast. This simplified depiction of current practice is a useful benchmark showing what the current practice approach would do without such ad hoc adjustments. The myopic dispatch without commitment policy is a natural benchmark; we will see that the current practice policy performs significantly better than this fully myopic policy.

We evaluate these policies against two performance bounds:

6. Perfect information (PI) dispatch. In every sample demand scenario, we find the minimum cost dispatch given the actual realized demands. This is equivalent to solving the UC problem with the actual demands for the given scenario rather than forecast demands.

7. Commitment with PI dispatch. Here, we use the commitments from the UC problem as in policy 4 above but then dispatch the uncommitted resources with perfect information about the realized demands.

Both of these PI policies are, of course, impossible to implement, as they rely on having perfect information about demands. PI dispatch provides a lower bound on the cost obtainable by any implementable policy; in our results (see, e.g., Table 1 in Section 6.4), we will compare the performance of all other policies against this PI benchmark. Considering the commitment-with-PI dispatch policy will help us disentangle the losses in the current practice approach (policy 4) because of committing units to the UC solution and losses from using myopic dispatch.

6.3. Simulation Details

We evaluate the seven policies described in Section 6.2 using Monte Carlo simulation, considering a total of 24

examples: eight days (one weekday and one weekend, drawn from each of the four seasons; see Table 1 for specific dates) evaluated in each of the three system scenarios (current, future, and future with carbon tax), as described above. Each of the 24 simulations uses 250 demand scenarios generated from a stochastic demand model, with the simulations starting at 6 a.m. on the given day and continuing for 24 hours.

6.3.1. Demand Model. As mentioned in Section 1, the net demand model that we use was developed by other members of the GRACE project team based on hourly forecast and actual demand data for 2019 provided by Duke Energy. The model uses a baseline forecast for net demand in each period and models errors from this baseline as evolving randomly according to an autoregressive process of high order. The data provided includes forecasts and actual demand as well as Duke-controlled solar supply; these solar supplies are netted out to generate forecasts and errors for net demand. The forecasts for the “future” system scenario were generated using this same model but with the solar forecasts and actuals multiplied by four. Figure 4 shows randomly generated net demand samples (light gray lines) for both the present and future systems for October 27. The solid black lines in the figure are the baseline forecasts. Comparing the two figures, we see that increasing the solar supply leads to a pronounced dip in afternoon net demand when solar production is high (as in Figure 1). The uncertainty in the afternoon is also amplified in this future system.

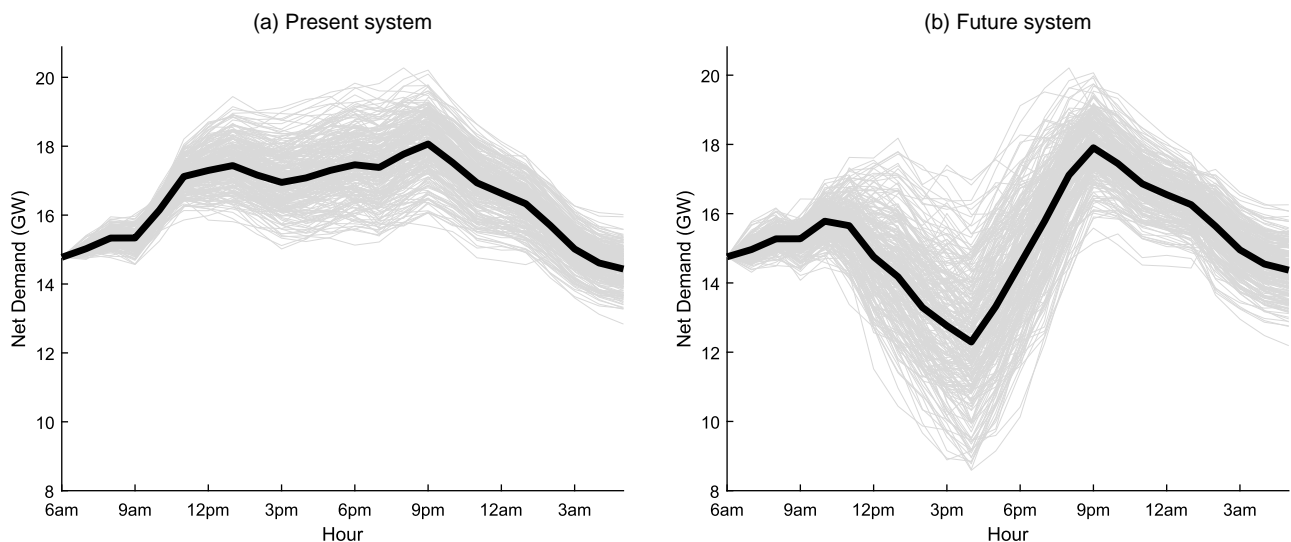
The state space associated with this high-order AR model is much too large to handle in the unit-specific DPs in the Lagrangian relaxation. In our numerical experiments, we use simple linear regression to fit an AR(1) model—with parameters varying by time of day and day—to simplify this high-order “real” demand

model. We then apply our weakly coupled DP approach and solve the resulting dual optimization problems (5) using a world-state process based on a discrete approximation of this AR(1) model involving 41 discrete demand states. The state spaces for the unit-specific DPs are thus 41 times the number of states included in the grid \mathcal{G}_s for the unit. As noted in Section 4, these grid sizes range from 1–10 for generation units (with an average of 2.60) and are 83 and 95 for the two storage units. With state and action spaces of this size, the unit-specific DPs can be quickly solved for any given price model.

In our simulations, we generate samples from the complex “real” demand model, and in the forward-looking ED problems (Section 5), we estimate the unit-specific expected continuation values using the value functions from this simplified demand model, with the values for specific demand levels given by interpolating these value functions between the discrete demand levels in the AR(1) approximation. Note that the Lagrangian performance bound of Proposition 1(b) applies in this simplified AR(1) model but does not provide a valid bound in this real model. Thus, in our numerical experiments, we use the costs associated with PI dispatch to provide performance bounds for this real model.

6.3.2. Initial Unit States. For each example day, we assume the day starts at 6 a.m. with an efficient set of units operating. Specifically, we sort the generation units by the full-power average costs (as in Figure 3) and assume the 6 a.m. demand is met by units with the cheapest possible full-power costs; all other units are assumed to be off. These assumed unit states are intended to represent realistic initial conditions for the system, given that demands at this time tend to be low. We assume both storage units are at 60% capacity, giving the system operator ample room to store or release energy, as desired, in the hours ahead.

Figure 4. Example October 27 Demand Scenarios



6.3.3. Terminal Values. Though our simulations are for a 24-hour period, the system operates over a very long time horizon, and we want our simulations to reflect costs over this longer horizon. For example, a policy that leaves storage empty at the end of the 24-hour period may achieve lower costs for these 24 hours than a policy that leaves storage full, but the costs of not storing energy would be borne in the days ahead. Similarly, a policy that leaves slow-start units with high start-up costs off might save costs in the short run but be costlier over a longer horizon. This is why, in current practice, one solves the UC problem with a one-week horizon to determine a baseline plan (and corresponding commitments) for the current day.

To reflect these longer-term considerations in the dynamic policies and simulations, for each day we simulate, we solve a weakly coupled DP (i.e., solve the Lagrangian dual (5)) for a one-week horizon to find

unit-specific value functions for this longer horizon; we use the simple price model for these long-horizon models. We then use period 25 of the 168-period value functions to serve as terminal values for the 24-hour model and the simulations; we subtract a constant from these value functions, so these terminal value adjustments will have zero mean under the policies that are optimal for these long-horizon DPs. These terminal values thus penalize policies that leave storage empty or are suboptimal in other ways for the longer horizon. We use the same penalties for all policies, including the PI policies.

6.4. Results

Table 1 summarizes the results of our numerical experiments, showing results for each of the eight days in each of the three system scenarios (present, future, and future with tax) for the policies discussed in Section 6.2. The values in the table are measures of the relative

Table 1. PI Gaps and Run Times

	Dynamic policies with price model			Commitment policies		Myopic
	Period constant	Period linear	Simple	Myopic dispatch	PI dispatch	Without commitment
January 27 (Sunday)						
Present	2.20%	0.14%	0.22%	1.45%	1.32%	3.30%
Future	3.46%	0.17%	0.20%	3.17%	2.83%	12.07%
Future plus carbon tax	3.32%	0.06%	0.09%	1.36%	1.12%	12.97%
February 6 (Wednesday)						
Present	2.42%	0.12%	0.18%	2.79%	2.49%	5.36%
Future	5.29%	0.45%	0.52%	6.89%	6.00%	25.75%
Future plus carbon tax	9.31%	0.74%	0.76%	6.42%	5.71%	39.51%
April 14 (Sunday)						
Present	3.32%	0.35%	0.37%	3.97%	3.64%	15.24%
Future	6.57%	0.50%	0.50%	12.99%	10.81%	32.27%
Future plus carbon tax	7.17%	0.49%	0.52%	9.98%	8.77%	44.46%
May 8 (Wednesday)						
Present	2.38%	0.30%	0.31%	2.35%	2.15%	12.38%
Future	4.94%	0.76%	0.87%	5.34%	4.47%	29.80%
Future plus carbon tax	5.34%	0.73%	0.85%	3.88%	3.36%	38.64%
July 7 (Sunday)						
Present	1.19%	0.09%	0.11%	0.98%	0.80%	12.91%
Future	2.31%	0.14%	0.19%	2.50%	2.07%	28.72%
Future plus carbon tax	2.08%	0.02%	0.03%	1.49%	1.25%	19.28%
July 31 (Wednesday)						
Present	0.97%	0.12%	0.16%	1.19%	0.91%	10.22%
Future	1.92%	0.11%	0.12%	2.55%	2.14%	25.30%
Future plus carbon tax	1.44%	0.01%	0.02%	0.99%	0.85%	18.42%
October 27 (Sunday)						
Present	3.14%	0.24%	0.26%	2.82%	2.62%	12.28%
Future	4.68%	0.28%	0.36%	8.44%	7.08%	28.85%
Future plus carbon tax	5.20%	0.31%	0.39%	6.53%	5.42%	39.51%
November 6 (Wednesday)						
Present	1.49%	0.09%	0.12%	1.64%	1.38%	12.33%
Future	3.17%	0.14%	0.20%	3.06%	2.42%	21.41%
Future plus carbon tax	1.71%	0.11%	0.24%	1.16%	0.84%	20.81%
Average gap	3.54%	0.27%	0.32%	3.91%	3.35%	21.74%
Present	2.14%	0.18%	0.22%	2.15%	1.91%	10.50%
Future	4.04%	0.32%	0.37%	5.62%	4.73%	25.52%
Future plus carbon tax	4.45%	0.31%	0.36%	3.98%	3.42%	29.20%
Average run time (seconds)	34.1	437.6	16.1	68.2		

performance of each policy compared with the unattainable performance bound of PI dispatch:

$$\text{PI gap} = \frac{(\text{average cost for policy}) - (\text{average cost for PI dispatch})}{(\text{average cost for PI dispatch})},$$

where the costs are averaged over the 250 simulated demand scenarios. Table EC-1 in the electronic companion provides mean standard errors (MSEs) for these PI gaps; these MSEs are small compared with the PI gaps shown in Table 1. Table EC-2 provides the PI gaps in absolute terms rather than relative terms.

The average run times for the policies are also shown in the table; all computations were done on a desktop computer, using MATLAB with the Gurobi optimization toolbox. For the dynamic policies, this run time is the time required to solve the Lagrangian dual problem (5) using a cutting-plane algorithm. The variation in run times reflects the number of parameters in the price model; models with more parameters tend to require more iterations in the cutting-plane algorithm (i.e., more cuts), with larger linear programs solved in each iteration and more work to compute gradients. For the current practice policy, the run time is the time required to solve the UC problem. The ED problems solved in each period in each demand scenario usually take a fraction of a second; we do not report these times.

6.4.1. An Example Day. Before discussing the results of Table 1, it is helpful to consider how the different policies behave on a single day. We focus on October 27 in the future system with carbon tax and consider one high- and one low-demand scenario. These extreme demand scenarios “stress test” the policies and make their differences clear.

Figures 5 and 6 show production decisions for these high- and low-demand scenarios. These figures are stacked area charts, where the stacked areas represent production from a particular unit; the units are stacked in the order of appearance in the supply curves of Figure 3. The colors in the figures correspond to those in the supply curves: blue denotes slow-start units, and pink denotes fast-start units. The thick black line represents the forecast demand, and in panels (b) to (d) of the figures, the thick blue line represents the realized (actual) demand in the given scenario. The light gray areas represent production from discharging storage units, and the dark gray areas represent charging the storage units.

We first consider the high-demand scenario in Figure 5. Figure 5(a) shows the UC plan designed to meet the deterministic forecast for this day. Here, the total production at 6 a.m. is slightly above the forecast demand, with the excess being stored. The UC plan uses storage to meet the morning and evening peaks (9 a.m. to 11

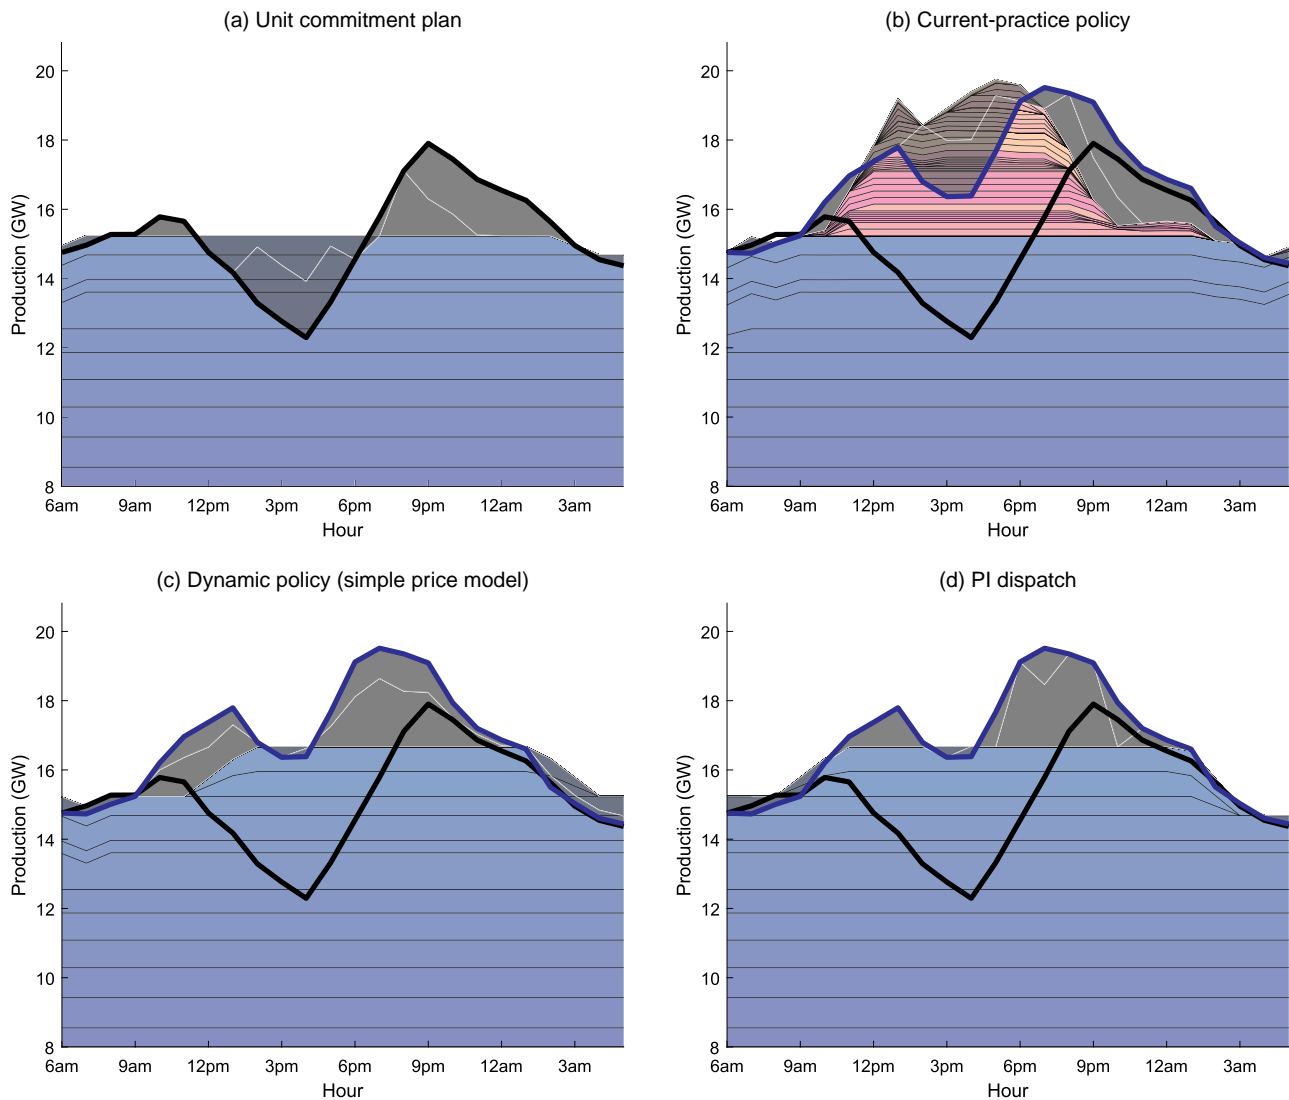
a.m. and after 6 p.m.) and stores energy through the afternoon. Given this UC plan, Figure 5(b) shows how the current practice policy operates in a high-demand scenario; in this scenario, the expected afternoon trough is replaced by a peak (imagine an unexpected cloud cover taking out the anticipated solar supply). Here, the midday peak and the commitment to store energy are covered using fast-start generation units. The costs of this policy are 12.4% higher than the costs of PI dispatch in this scenario. In contrast, the dynamic dispatch policy (with the simple price model), shown in Figure 5(c), meets these midday peaks through a combination of starting additional slow-start units (around 11 a.m.) and increased use of stored energy. This plan is quite similar to PI dispatch (see Figure 5(d)), and the PI gap is only 0.3%. PI dispatch does slightly better than the dynamic policy because, with perfect information about the afternoon peak, it starts the additional slow-start units a bit earlier and optimizes the use of storage accordingly.

Figure 6 shows similar plots for a low-demand scenario. In this low-demand scenario, the PI gap for the current practice policy is 16.2%, compared with only 0.6% for dynamic dispatch. As shown in Figure 6, (c) and (d), the dynamic and PI policies shut down some of the operating slow-start units heading into the deeper-than-expected afternoon trough and store more through this trough. Instead of shutting down the slow-start units, the current practice policy is committed to keeping these units on and curtails production at many of these slow-start units while continuing to incur their fixed costs. Being committed to their storage plan, the current practice policy sheds a significant amount of supply—shown in red—from 3 p.m. to 6 p.m., despite the fact that this energy could be stored; a real-life system operator would likely adjust the storage plan in this situation.

Figure 7 shows costs for the dynamic and current practice policies for all 250 scenarios for October 27 for the present system as well as the future system with carbon tax. Here, we see that the dynamic policies outperform the current practice policy in *every* demand scenario, not just the extreme scenarios discussed earlier. As noted in Table 1, the average PI gaps are 0.26% and 2.82% for the dynamic and current practice policies, respectively, in the present system (corresponding to an average improvement of $\approx \$107,000$) and 0.39% and 6.53% in the future system with carbon taxes (an improvement of $\approx \$605,000$). In the future system without carbon taxes (not shown in Figure 7), the PI gaps are 0.36% and 8.44% for the dynamic and current practice policies, respectively (an improvement of $\approx \$230,000$).

6.4.2. All Day and System Scenarios. The results for the other days considered are similar to those for this example day. The best-performing policy is always the dynamic policy using the period-linear price model,

Figure 5. (Color online) Production Plans for October 27 in a High-Demand Scenario

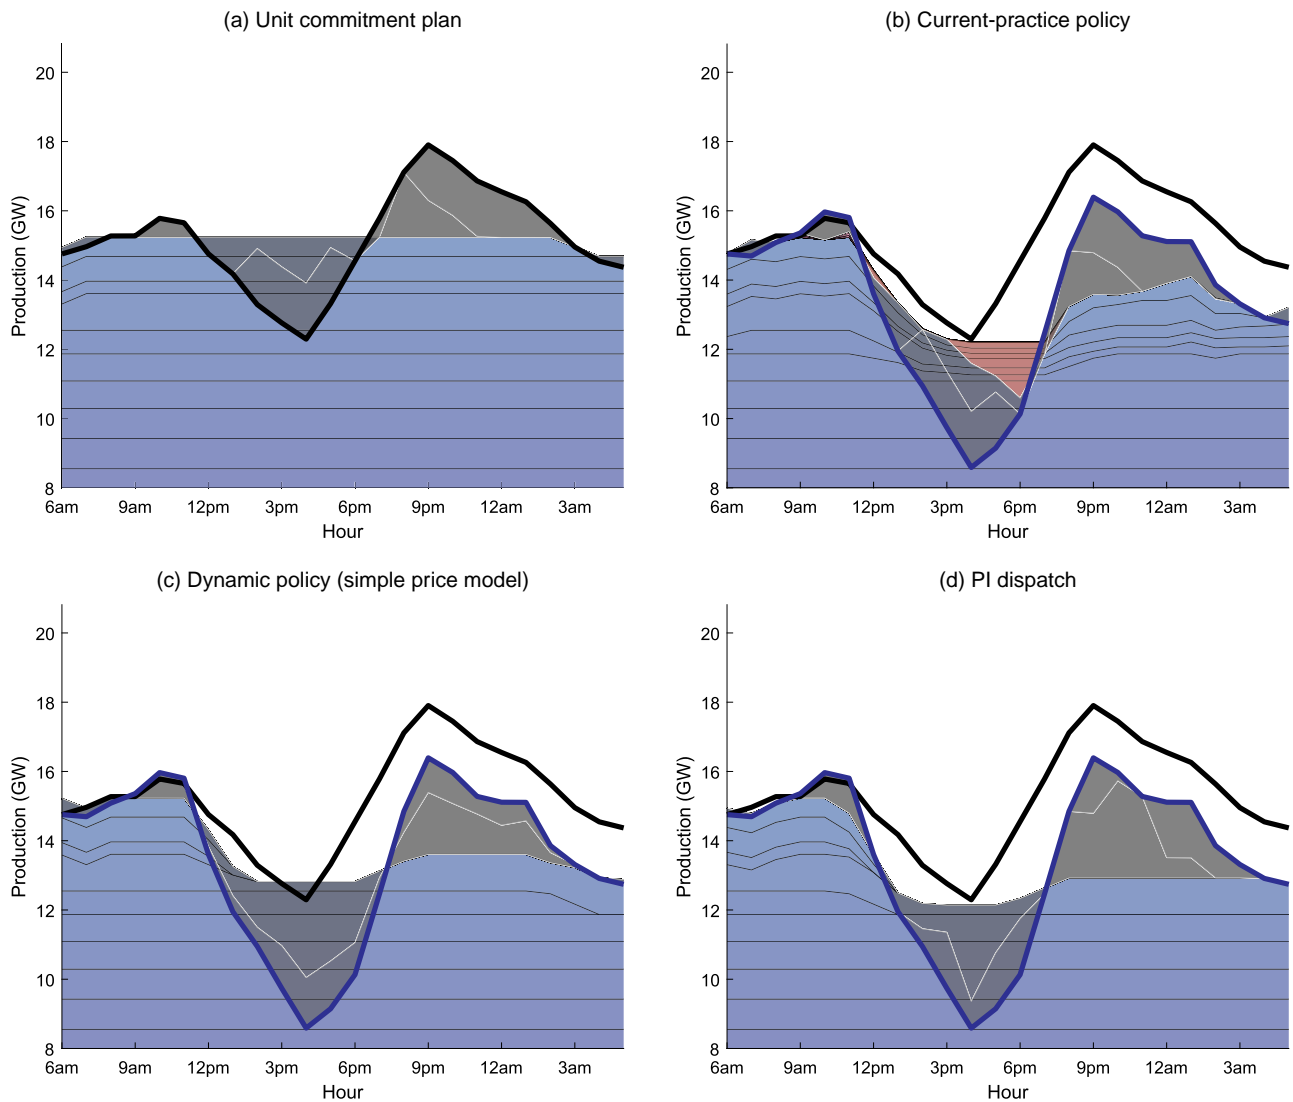


followed closely by the dynamic policy with the simple price model. The period-constant price model performs significantly worse than the two stochastic price models, highlighting the importance of having prices vary with the uncertainty in demand. The current practice policy (commitment with myopic dispatch) performs worse than these two dynamic dispatch policies every day and in every system scenario.

It is natural to wonder whether the relatively poor performance of current practice is because of its reliance on commitments or because of the use of myopic dispatch. The results for the commitment policy with PI dispatch suggest that the issue is primarily the commitments: even with perfect information in dispatch, the performance does not improve much compared with commitment with myopic dispatch. Nonetheless, the current practice policy is much better than a truly

myopic policy: when paired with myopic dispatch, commitment serves as a way to force adherence to a longer-term plan.

Finally, we note that differences between the proposed dynamic and current practice policies are more pronounced in the future system scenarios. Averaging across the eight days considered, the proposed dynamic price model reduces expected costs by $\approx \$101,000$ in the present system and $\approx \$188,000$ in the future system. With the additional solar (four times) and storage (two times), the uncertainty in supply and the flexibility in storage both increase, and the benefits of the dynamic approach—handling uncertainty explicitly and better using the system’s flexibility—also increase. In the future system, the addition of the carbon tax reduces the percentage improvement from current practice to dynamic policies. As noted in the discussion of the system supply

Figure 6. (Color online) Production Plans for October 27 in a Low-Demand Scenario

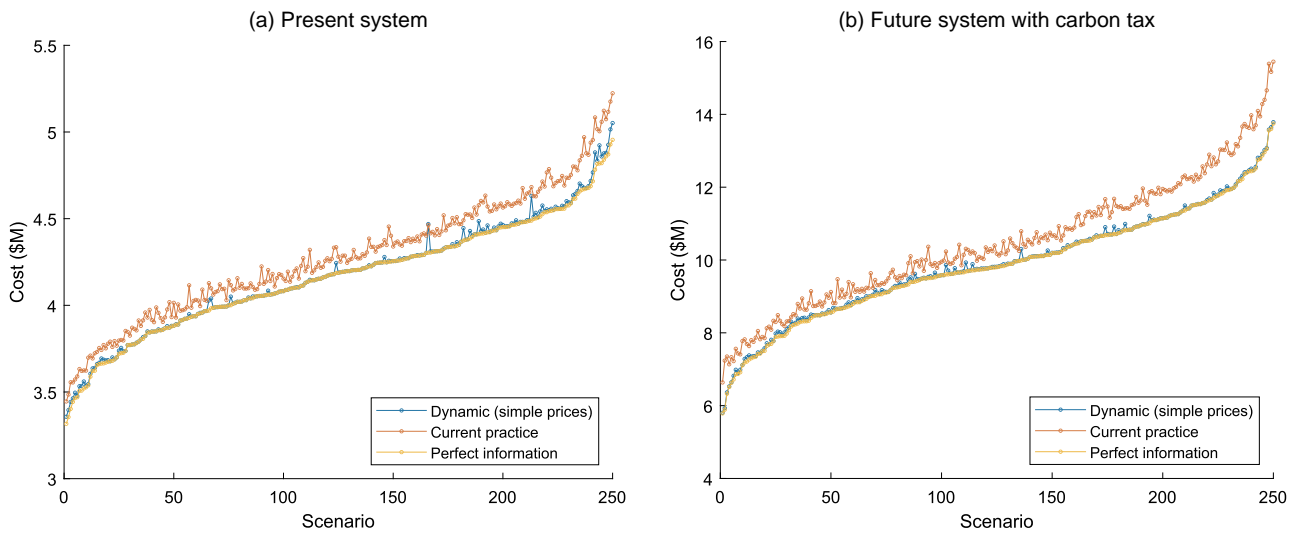
curves in Section 6.1, the carbon tax effectively eliminates the slow-start, coal-fired generation units from the supply mix. These slow-start units are better managed by the dynamic policies than the current practice policies; removing the coal-fired slow-start units somewhat neutralizes this advantage. However, because the costs are so much higher with the carbon tax, in absolute terms (e.g., in dollars rather than percentages), the improvement in going from current practice to dynamic policies is greater when carbon taxes are included in this future scenario, with an average improvement of $\approx \$413,000$.

6.5. Unit Values

In addition to providing the basis for dynamic dispatch, the Unit-Specific Value Functions (4) of the Lagrangian model are interesting in their own right. Recall that

these value functions can be interpreted as the expected total profit generated by the unit over the 24-hour-period, taking into account the terminal values for the ending state. Figure 8 shows the value-per-unit capacity for the generation units in the model system for October 27 (the day featured in Sections 6.3 and 6.4) and July 31, both in the future system with the carbon tax scenario. The values plotted are the value of the unit given the unit's initial state and the initial world-state, divided by the capacity of the unit. The widths of the plotted lines for each unit are the unit's capacity, and the units are sorted and accumulated as in the supply curves of Figure 3, using the same colors for slow-start and fast-start units (blue to green and pink to yellow, respectively). Sample demand forecasts for July 31 are shown in Figure EC-1 in the electronic companion; this day has

Figure 7. (Color online) Costs for 250 Demand Scenarios for October 27



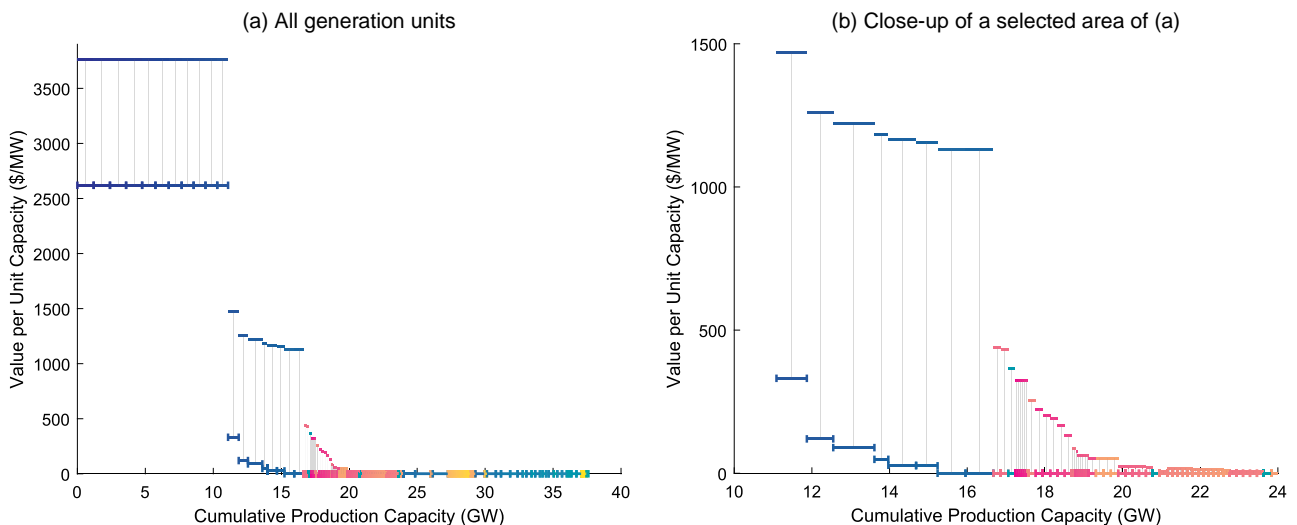
Note. Scenarios are sorted by the cost of PI dispatch.

a significant evening peak, with the forecast mean at 9 p.m. being ≈ 26 gigawatts, but less uncertainty than October 27.

Comparing the unit values for these two days, we see that the unit values for October 27 (with hatched lines) are much lower than those for July 31 (with flat lines); this is because prices are generally lower on October 27. On both days, the nuclear units are the most valuable, as they provide free energy with no carbon emissions; the natural gas slow-start units are the next most valuable. On October 27, none of the fast-start units have any value; this is because, even in the

scenarios with the highest realized demand (see Figure 5(c)), the peaks can be covered with storage. In contrast, on July 31, some of the natural gas fast-start units have positive values because these units *may* be deployed (along with storage) to get through the peak. These unit values, based on a stochastic model, reflect these possibilities. On both days, many of the fast-start units and all of the coal-fired units have zero value because they are never deployed. The two storage units are worth more on October 27 than on July 31 (7% and 14% more), reflecting the greater uncertainty on October 27.

Figure 8. (Color online) Unit Values for October 27 and July 31



Note. Future system with carbon tax.

7. Conclusions

This paper develops a dynamic approach based on weakly coupled DP for managing production under uncertainty in an integrated energy system. An essential component of this approach is choosing stochastic prices that lead to unit-specific DPs that are easy to solve while simultaneously providing good performance. We are encouraged by the fact that “simple” price and world-state models are effective in our numerical examples despite the complexity of the demand models and the size of these systems. The proposed approach scales gracefully with the number of units, which suggests this approach is viable in even larger systems.

Several extensions may be relevant in practice. First, generation units occasionally fail because of weather events or idiosyncratic reasons. We could augment the state spaces for the units to include outage states or, more generally, “derating” states where units have reduced production. Second, though this is not a major issue in the current Duke Energy system, in many other systems, transmission constraints restrict distribution. As discussed in Section EC-4 in the electronic companion, we can relax demand-matching constraints separately at different nodes in the network, with prices varying across nodes. These variations may complicate the unit-specific DPs and how we choose basis functions for the price model. It would be helpful to study some examples of such network systems in detail. Finally, we have presented example models of generation and storage units; one could develop other unit-specific models for assets with other features. We leave a detailed study of these extensions for future research.

We have focused on a vertically integrated utility that manages its own generation assets. Such vertically integrated utilities provide about one-third of the electricity generated in the United States, with the remaining two-thirds of generation managed by independent system operators (ISOs) working in a market setting. In these markets, generators provide bids to the ISO that describe their generation units, including start-up costs, no-load costs, fuel costs, etc., as discussed in Section 4.1. Given these bids, in the day-ahead market, ISOs solve a version of the UC problem to determine a deterministic operating plan for the next day. Then, in the real-time markets, the ISOs solve an ED problem (typically every five minutes) to ensure that production meets actual demand. The methods described in this paper would thus potentially apply to ISOs as well as vertically integrated utilities.

However, the ISO setting introduces additional complexities. Notably, ISOs use the shadow prices associated with the demand constraints (and network constraints) in the UC problem to settle the market: consumers of energy pay, and producers are compensated based on the prices associated with this deterministic

plan. The real-time markets similarly settle based on shadow prices from the ED problem. Our Lagrangian relaxation of the DP (Section 3) serves a role analogous to the UC problem in terms of developing plans for the coming day. However, this stochastic model considers many possible scenarios and prices, and there is no obvious method for settling the day-ahead market. Our forward-looking ED problem (Section 5) is a deterministic optimization problem that can provide shadow prices, but these shadow prices reflect the longer-term values included in the objective function (see equation (20)) and thus differ from those used by ISOs in real-time markets. Developing appropriate market mechanisms to accommodate the increasing uncertainty and flexibility in energy markets is an important area for future research.

In future work, we would also like to study the duality gaps—differences between Lagrangian bounds and the value for the original DP (as in Theorem 2)—when there are constraints on the price model, and we characterize the forward-looking ED-based dispatch policy’s performance with respect to these bounds. Finally, we would like to better understand the principles of selecting good basis functions for price models in applications.

Acknowledgments

The authors thank Antonio Conejo, David Hsieh, Erin Mansur, David McAdams, Dalia Patiño-Echeverri, and seminar participants at Fuqua, Tuck, the University of Toronto, INSEAD, and INFORMS conferences, as well as the area editor, an anonymous associate editor, and two anonymous referees for their helpful comments. The authors also gratefully acknowledge GRACE teammates Dimitrios Floros and Jordan Kern for providing the net demand forecasting models that we use in our numerical examples and Duke Energy for providing the data underlying these examples and insight into the unit commitment problem. The processes and procedures described in this paper, as well as any products derived from its content, are the subject of a pending patent owned by Duke University and Dartmouth College, who also own the copyrights in the software implementing the concepts. All rights are reserved.

Endnotes

¹ One may wonder why storage is committed to the UC solution, given the inherent flexibility of storage resources. The ED problems are typically formulated myopically, minimizing the cost for a given period. With such a myopic approach, the ED solutions would frequently draw down storage rather than save stored energy for a later peak, as may be prescribed in the UC solution. We will see in Section 6.4 that this current practice approach with commitment indeed performs significantly better than a fully myopic approach.

² The assumption that $V_{T+1}^*(x, \psi_{T+1}) = 0$ is not critical. The Lagrangian decomposes across units (see Section 3) if the terminal value function is additively separable across units. For the structural properties of the unit value functions (see Section 4) to hold, the terminal value functions must also satisfy these structural properties.

In our numerical work (see Section 6.3), we will use terminal value functions that are Lagrangian value functions for a model with a longer time horizon that has zero terminal values. These terminal value functions thus decompose and have the desired structure.

³ The assumption of Theorem 2 that units' production possibilities are convex is required to evaluate unit costs for "mixed" actions. Though this condition is satisfied for storage units, generation units may have nonconvex production possibilities. For example, using the notation of Section 4, the possible production levels for a generation unit may be $\{0\} \cup [p_s, \bar{p}_s]$, where zero corresponds to an off state with zero production, and $\bar{p}_s \geq p_s > 0$. Thus, one cannot achieve production levels between zero and p_s . We can, however, convexify these production possibility sets by augmenting generation units to allow them to shed production, albeit at some cost: one can achieve levels between zero and p_s by producing at p_s and shedding the excess production.

⁴ If there are minimum downtime constraints requiring a unit to be off for at least ℓ periods when shut down, we could augment the unit state space to indicate how long a unit has been off. Alternatively, we can replace the $t + 1$ time index in the continuation value when shutting down in (4) with $t + \ell$ so the unit is not available to be started up again until period $t + \ell$. Minimum uptime constraints can be handled similarly; see Takriti et al. (1996).

⁵ Our model assumes that the costs when producing are linear in the production level. In practice, it is sometimes assumed that these costs are piecewise linear in production. The result of Proposition 4 also holds with piecewise linear convex costs provided we augment the set of grid points to include the breakpoints of the piecewise linear cost function in the set of possible values \hat{x} in (18) the generate the grid. (Convexity of the piecewise linear cost functions is only needed for the concavity claim in Proposition 4(a).)

⁶ This is a 24-period deterministic example; we randomly generated prices to create an "interesting" value function. In this particular scenario, a low price is expected in two periods: the sharp drop in the value function at 0.5 gigawatts reflects the fact that for production states above this level, it is not possible to ramp down enough (to $p_{s,d} = 0.3$ or less) to shut the unit down before the low price period.

⁷ As written, this MILP allows the solution to interpolate between nonadjacent grid points. We could improve this formulation by introducing additional binary variables and constraints that require at most two consecutive weights to be nonzero; see, for example, Bertsimas and Tsitsiklis (1997, p. 455). However, this could be expensive from a computational perspective, as it requires introducing an additional set of binary variables (one for each feasible action) for each generation unit.

⁸ Note that we are not explicitly modeling the utility-scale solar as being curtailable, though it may be curtailed in practice. In our examples, we never shed supply with our dynamic policies, and supply is rarely shed with other policies. One could introduce additional units representing the curtailable production. Doing this would require explicitly modeling curtailable solar supply in addition to net demand. In the Lagrangian relaxation, the corresponding unit DPs would produce all of the available solar supply whenever the price is positive, and zero otherwise. In dispatch, the ED solutions may choose levels between these extreme solutions and produce some of the available solar.

References

Adelman D, Mersereau AJ (2008) Relaxations of weakly coupled stochastic dynamic programs. *Oper. Res.* 56(3):712–727.
Barty K, Carpentier P, Girardeau P (2010a) Decomposition of large-scale stochastic optimal control problems. *RAIRO-Oper. Res.* 44(3):167–183.

Barty K, Carpentier P, Cohen G, Girardeau P (2010b) Price decomposition in large-scale stochastic optimal control. Preprint, submitted December 9, <https://arxiv.org/abs/1012.2092>.
Bertsekas DP (1996) *Constrained Optimization and Lagrange Multiplier Methods* (Athena Scientific, Belmont, MA).
Bertsekas DP, Nedić A, Ozdaglar AE (2003) *Convex Analysis and Optimization* (Athena Scientific, Belmont, MA).
Bertsimas D, Mersereau AJ (2007) A learning approach for interactive marketing to a customer segment. *Oper. Res.* 55(6):1120–1135.
Bertsimas D, Tsitsiklis JN (1997) *Introduction to Linear Optimization*, Athena Scientific Series in Optimization and Neural Computation, vol. 6 (Athena Scientific, Belmont, MA).
Brown DB, Smith JE (2020) Index policies and performance bounds for dynamic selection problems. *Management Sci.* 66(7):3029–3050.
Brown DB, Zhang J (2022) Dynamic programs with shared resources and signals: Dynamic fluid policies and asymptotic optimality. *Oper. Res.* 70(5):3015–3033.
Brown DB, Smith JE, Sun P (2010) Information relaxations and duality in stochastic dynamic programs. *Oper. Res.* 58(4-part-1):785–801.
Caro F, Gallien J (2007) Dynamic assortment with demand learning for seasonal consumer goods. *Management Sci.* 53(2):276–292.
Carpentier P, Gohén G, Culioli J-C, Renaud A (1996) Stochastic optimization of unit commitment: A new decomposition framework. *IEEE Trans. Power Systems* 11(2):1067–1073.
Cerisola S, Baíllo Á, Fernández-López JM, Ramos A, Gollmer R (2009) Stochastic power generation unit commitment in electricity markets: A novel formulation and a comparison of solution methods. *Oper. Res.* 57(1):32–46.
Conejo AJ, Baringo L (2018) Unit commitment and economic dispatch. *Power System Operations. Power Electronics and Power Systems* (Springer International Publishing, Cham, Switzerland), 197–232.
Dentcheva D, Römisch W (1998) Optimal power generation under uncertainty via stochastic programming. Marti K, Kall P, eds. *Stochastic Programming Methods Tech. Appl.*, Lecture Notes in Economics and Mathematical Systems, vol. 458 (Springer, Berlin, Heidelberg), 22–56.
Duke Energy (2023) 2023 Carolinas resource plan. Accessed September 23, 2023, <https://www.duke-energy.com/our-company/about-us/irp-carolinas>.
Hawkins JT (2003) A Lagrangian decomposition approach to weakly coupled dynamic optimization problems and its applications. PhD thesis, Massachusetts Institute of Technology, Cambridge.
Khazaei J, Powell WB (2018) Smart-invest: A stochastic, dynamic planning for optimizing investments in wind, solar, and storage in the presence of fossil fuels: the case of the PJM electricity market. *Energy Systems* 9:277–303.
Miao S, Jasin S, Chao X (2022) Asymptotically optimal Lagrangian policies for multi-warehouse, multi-store systems with lost sales. *Oper. Res.* 70(1):141–159.
Nowak MP, Römisch W (2000) Stochastic Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty. *Ann. Oper. Res.* 100:251–272.
Nürnberg R, Römisch W (2002) A two-stage planning model for power scheduling in a hydro-thermal system under uncertainty. *Optim. Engrg.* 3:355–378.
Patino-Echeverri D, Principal Investigator (2023) GRACE: A grid that is risk aware for clean electricity. Accessed September 23, 2023, <https://arpa-e.energy.gov/programs-and-initiatives/search-all-projects/grid-thats-risk-aware-clean-electricity-grace>.
Powell WB, Meisel S (2015) Tutorial on stochastic optimization in energy—Part I: Modeling and policies. *IEEE Trans. Power Systems* 31(2):1459–1467.
Rennert K, Errickson F, Prest BC, Rennels L, Newell RG, Pizer W, Kingdon C, et al. (2022) Comprehensive evidence implies a higher social cost of CO₂. *Nature* 610(7933):687–692.

- Tahanan M, Van Ackooij W, Frangioni A, Lacalandra F (2015) Large-scale unit commitment under uncertainty. *4OR-Quart. J. Oper. Res.* 13:115–171.
- Takriti S, Birge JR (2000) Lagrangian solution techniques and bounds for loosely coupled mixed-integer stochastic programs. *Oper. Res.* 48(1):91–98.
- Takriti S, Birge JR, Long E (1996) A stochastic model for the unit commitment problem. *IEEE Trans. Power Systems* 11(3): 1497–1508.
- Topaloglu H (2009) Using Lagrangian relaxation to compute capacity-dependent bid prices in network revenue management. *Oper. Res.* 57(3):637–649.
- Whittle P (1988) Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.* 25(A):287–298.

David B. Brown is the Snow Family Business Professor and the Faculty Director for the Center for Energy, Development and Global Environment (EDGE) at Duke University's Fuqua School of Business. In his research, he develops and analyzes methods for stochastic optimization problems, including applications in energy, finance, and operations management.

James E. Smith is the Jack Byrne Distinguished Professor in Decision Science at the Tuck School of Business at Dartmouth College. His research interests focus on developing methods for dynamic decisions under uncertainty. He is part of Dartmouth's Jack Byrne Academic Cluster in Mathematics and Decision Science, which focuses on developing and applying mathematical thinking to societal challenges in fields such as healthcare, energy, and manufacturing and in understanding the impacts of climate change.