Energy Systems manuscript No.
(will be inserted by the editor)

A Stochastic Dynamic Programming Model for Co-optimization of Distributed Energy Storage

Xiaomin Xi · Ramteen Sioshansi · Vincenzo Marano

the date of receipt and acceptance should be inserted later

Abstract We develop a stochastic dynamic programming model that co-optimizes the use of energy storage for multiple applications, such as energy, capacity, and backup services, while accounting for market and system uncertainty. Using the example of a battery that has been installed in a home as a distributed storage device, we demonstrate the ability of the model to co-optimize services that 'compete' for the capacity of the battery. We also show that these multiple uses of a battery can provide substantive value.

Keywords Energy storage · stochastic dynamic program · co-optimization

1 Introduction

Developments in the electricity industry over the past few years have increased interest in energy storage. One of the reasons for this interest is the introduction of markets that signal the cost and value of the services that storage can provide. Another is the increasing strains that renewable energy is placing on electric power systems and the role that storage can play in mitigating these integration issues.

A number of papers study potential uses of storage and estimate the value of some of those applications based on historical data. EPRI [10] provides one of the first discussions of storage. Because the discussion is framed in the 1970s, before the advent of restructured electricity markets, it focuses on storage use

X. Xi · R. Sioshansi
Integrated Systems Engineering Department, The Ohio State University, Columbus, Ohio, United States
Phone: +1-614-292-3932
E-mail: xi.12@osu.edu, sioshansi.1@osu.edu

V. Marano
Center for Automotive Research, The Ohio State University, Columbus, Ohio, United States
E-mail: marano.8@osu.edu

by a vertically integrated utility to avert the need for peaking generation capacity by shifting on-peak loads to off-peak periods. More recent discussions of storage focus on the role of markets [11,13]. These papers provide comprehensive overviews of storage technologies and uses, including energy, capacity, renewable-related, transmission, distribution, and consumer applications.

One of the most commonly studied storage applications is what is often called 'energy arbitrage'—charging storage when wholesale energy prices are low and discharging when high to take advantage of diurnal price differences [16,3,14]. Many of these analyses use a perfect-foresight assumption, under which prices are assumed to be known perfectly when making storage decisions. This assumption allows the value of arbitrage (and other) services to be estimated using historical price data and patterns.

Other analyses expand these works by relaxing the perfect-foresight assumption or by considering arbitrage in conjunction with other uses. Mokrian and Stephen [24] describe a stochastic dynamic programming (SDP) model to maximize expected arbitrage revenues while accounting for energy price uncertainty. Sioshansi *et al.* [35,34] relax the perfect-foresight assumption by examining a 'backcasting' heuristic whereby storage is dispatched using historical price patterns that are assumed to repeat themselves. Walawalkar *et al.* [39] examine storage economics in the New York ISO market, considering both arbitrage and ancillary services (AS). AS are excess generating capacity that a utility or system operator (SO) reserves in order to provide a buffer for real-time deviations between actual and forecasted demand or supply of energy. They find that storage has a high probability of yielding a positive net present value if installed in New York City. Drury *et al.* [9] examine the value of arbitrage and AS from storage in a number of markets in the United States. Other papers examine interactions between storage and renewables. This includes the use of storage to mitigate renewable variability and uncertainty [26, 1,15]; the economic and emissions impacts of storage and renewables [6,17,32, 31]; and using storage to reduce the need for dedicated transmission to deliver renewable energy to load centers [7,33].

There are, however, numerous issues and nuances of storage that are not well addressed by this literature. One is that most analyses consider only a single storage application and do not co-optimize multiple storage uses. Evaluating multiple uses of storage is critically important, given the high capital costs of most storage technologies. Most storage analyses that consider only one application find that storage is not economic on the basis of that single use. Since different storage applications can conflict with each other, determining the value of multiple storage applications requires those applications to be co-optimized. For instance, if storage is discharged at time t to earn arbitrage revenue this can conflict with its ability to provide AS at time $s \geq t$, since there may be less energy in storage in the future. Although Walawalkar *et al.* [39] analyze the value of AS and arbitrage, they do not employ an optimization method in their analysis. Rather, they rely on heuristics based on price duration curves and patterns. Similarly, the analysis of Drury *et al.* [9] is lim-

ited in that they do not fully account for the uncertain interactions between providing energy and AS.

The effects of price and system uncertainty are also often neglected in storage analyses. Price uncertainty can impact storage use, since charging, discharging, and other decisions may not be value-maximizing depending on the stream of realized prices. Although Mokrian and Stephen [24] develop a model that maximizes expected arbitrage revenue while accounting for price uncertainty, this neglects other forms of uncertainty that can also impact storage. For instance, if storage provides AS at time t, the amount of energy in storage at time $s > t$ varies depending on how much energy must actually be provided as a result of the reserved capacity being called. This can affect the ability of storage to provide other services in the future.

Another important facet of storage is the potential scale of the technology. Many storage analyses consider 'utility-scale' storage—devices that can charge and discharge hundreds of MW for multiple (in some cases more than 20) hours. Although most storage is currently utility-scale,[1] smaller-scale storage is also becoming an attractive option, due to the ability of such distributed storage to provide services that utility-scale storage cannot. For example, American Electric Power (AEP) installed a 1 MW sodium-sulfur (NaS) battery facility in a community in West Virginia to relieve a constraint on a distribution-level transformer [25].

Given these limitations of existing storage modeling techniques, this paper describes an SDP model that co-optimizes multiple storage applications while accounting for market and system uncertainty. Using this model we study the value of a battery that is installed in a residential home as a stationary storage device. We use a discretization-based approach to derive near-optimal policies and provide bounds on the optimality gap. This analysis demonstrates how the operation and value of the battery varies depending on the combination of applications considered. Although our case study focuses on small-scale distributed storage, our model is sufficiently general that it can be applied to other cases, including utility-scale storage.

The remainder of this paper is organized as follows. Section 2 describes in further detail services that storage, particularly distributed storage, can provide. This discussion also motivates the storage applications that we focus on in our analysis. Section 3 provides the formulation of our model and Section 4 discusses the discretization technique that we use to find near-optimal solutions. Section 5 discusses the assumptions and data underlying the case study that we examine. Section 6 summarizes our results, and Section 7 concludes.

[1] Pumped hydroelectric storage (PHS) accounts for most currently installed storage capacity [5] and interest in compressed-air energy storage (CAES) is increasing [36]. Due to their physical attributes and geological requirements, PHS and CAES are both utility-scale storage technologies.

2 Uses of Distributed Storage

Our model assumes distributed storage, whereby batteries are installed in residential homes. We do not make an explicit assumption regarding battery ownership—they could in principle be owned by the homeowner, utility, or an aggregator that provides storage using batteries deployed in multiple homes. We assume that the decision maker that optimizes battery use (which can be a different entity than the battery owner) is exposed to wholesale market price signals, for instance real-time energy and AS prices. The homeowner does not have to be exposed to or pay such prices, however. For example, the batteries could be owned by the utility, which optimizes their use against real-time prices, while the homeowner pays a fixed time-invariant retail electricity rate. We assume that the battery pack includes an automated control system that charges and discharges the battery, or a communication system that allows centralized control by the decision maker. Although they are not currently ubiquitous, such systems have existed for some time. Examples include the automated air conditioning cycling system implemented by Southern California Edison beginning in 1980 [12], and the communication and control systems used by the NaS battery deployed by AEP in West Virginia. Given these assumptions utility or aggregator ownership may be the easiest contractually, although other ownership structures could be used. For example, the homeowner could own the battery and lease it to the utility, which provides the homeowner with an electricity bill rebate. The battery pack would include a communication system that allows the utility to operate it remotely. The homeowner continues to pay a fixed retail rate for its electricity, while the utility optimizes battery charging and discharging.

Although storage can provide many services, we focus on four that are particularly conducive to this case: energy arbitrage, AS, backup energy, and relief of distribution constraints. We discuss in Section 3 how our model can capture other uses that we do not consider here. An entity that sells capacity in the AS market is contractually obligated to provide energy in real-time if called upon by the SO. AS are typically given two-part payments. One is a capacity payment, which is paid regardless of whether the reserved capacity is actually needed in real-time. The other is an energy payment, which is paid if energy is called in real-time. An entity that sells AS capacity but cannot provide called energy is penalized for its shortfall. The penalty rate is typically a fixed percentage premium over the real-time energy price. AS are categorized into different quality levels, depending on the rate at which the provider is obligated to respond to a call for energy in real-time. Three common types of AS are regulation, and spinning and non-spinning reserves. Regulation is the highest-quality of all AS, which is used to fine-tune the frequency and voltage of the grid by exactly matching real-time energy demand and supply. Thus, regulation requires the greatest flexibility on the part of the supplier. Regulation is further subdivided into regulation-up and -down services. A provider of regulation up must increase its output from a baseline level if called in real-time, whereas a regulation down provider must decrease it. Some SOs

treat regulation up and down as two separate capacity products, and in such a market an entity could sell different amounts of the two capacities. Others treat them as a single product, and an entity that sells regulation capacity in such a market is obligated to provide both regulation-up and -down energy, if required. Although energy storage can provide different AS, regulation is significantly more valuable than other AS [37]. Thus we focus on regulation services only in our analysis.

Backup energy refers to using the battery as an energy source for the home in the event of a supply disruption, such as a generation, transmission, or distribution outage. This storage application could be of great value to both the end user and the utility. Electricity service disruptions are inconvenient and costly to end users and customers derive value from averting an outage. For instance, the cost of the blackout that affected the northeastern United States and Canadian province of Ontario in 2003 is estimated to be between $6.4 billion and $10 billion [18]. A utility could also derive value since the battery could help it meet reliability standards. This is one of the ancillary benefits of the NaS batteries that AEP installed in West Virginia [25]. Even if a battery does not have sufficient storage capacity to provide energy throughout a prolonged outage, a limited energy supply could reduce the severity of an outage.

The final application that we consider is the use of batteries to relieve distribution constraints. Increasing use of electronic devices can create bottlenecks and strains on distribution systems. The introduction of plug-in electric vehicles (PEVs), which can require high-power fast charging, can exacerbate this issue. Mohseni and Stevie [23] examine PEV integration in two regions of the United States and predict significant risks to utilities due to potential geographic clustering of PEV owners. For instance, if several PEVs in a neighborhood begin charging in the early evening when they arrive home, the total vehicle and building loads can overload a distribution-level transformer. Batteries installed in-home can relieve these strains by shifting loads to periods with lower distribution-level loads. This can also be useful without PEVs, since increasing use of electric devices can overload the electric circuit leading into a home, which a battery could relieve.

3 Model Formulation

With the applications that we study, storage decisions depend on two markets, energy and regulation, the state of the system (*i.e.*, whether there is an outage or not), and an exogenous building energy demand. In each time period, the storage operator makes battery charging, discharging, and regulation sales decisions. If the system is not in an outage state, the battery can be charged or discharged from the grid and can provide regulation. Otherwise, the battery can only be discharged to serve the building load. We assume in our analysis that all decisions are made at hourly timesteps, although this can be relaxed by appropriately scaling the parameters and variables. The model assumes that

the battery is sufficiently small compared to the total market that storage decisions do not impact prices or otherwise affect the system.

Our model formulation follows the conventions and notation that Powell [27] uses. We give the formulation by first defining the relevant parameters and variables, and then giving state-transition and objective functions. We use the convention that variables with a subscript t are unknown (stochastic) before hour t and become known (deterministic) at hour t.

3.1 Parameters

$\overline{P}^b$: maximum charging and discharging power capacity of the battery [kW]
$\overline{R}$: maximum storage level of the battery [kWh]
$\underline{R}$: minimum storage level of the battery [kWh]
η^c: charging efficiency of the battery
η^d: discharging efficiency of the battery
$\overline{P}^h$: maximum power capacity of the home [kW]
V^L: penalty for unserved building energy [$/kWh]
V^R: penalty for unserved regulation call
γ: discount factor

The power capacity of the battery, $\overline{P}^b$, typically reflects power limits of the inverter in the battery pack or thermal limits of the electrochemical storage medium of the battery (batteries generate heat when charged and discharged). The minimum storage level, $\underline{R}$, arises because some battery technologies suffer extreme cycle-life degradation if the state of charge (SOC) falls too low. For instance, lithium-ion batteries can typically only cycle down to a 30% SOC, and this type of a restriction is included in our model. The unitless ratios, η^c and η^d, reflect efficiency losses from battery charging and discharging. The product of these two terms, $\eta^c \cdot \eta^d$, gives the roundtrip efficiency of the battery (*i.e.*, kWh of energy that can be discharged in net per kWh of energy charged into the battery). $\overline{P}^h$ is the total power capacity of the building circuit. V^L is the cost associated with unserved building load. V^R captures the penalty imposed by the SO for an unserved regulation call, which is modeled as being a fixed percentage of the prevailing wholesale energy price.

3.2 Decision (Action) Variables

e_t^d: energy discharged from battery in hour t for energy sales [kWh]
e_t^c: energy charged into battery in hour t [kWh]
e_t^l: energy discharged from battery in hour t to serve building load [kWh]
l_t: building load met in hour t [kWh]
k_t^u: regulation-up capacity sold in hour t [kW-h]
k_t^d: regulation-down capacity sold in hour t [kW-h]

We also define $a_t = (e_t^d, e_t^c, e_t^l, l_t, k_t^u, k_t^d)$ as a vector of hour-t decision variables.

3.3 State Variables

x_t: total energy in storage at the beginning of hour t [kWh]
p_t^e: hour-t market price of energy [$/kWh]
p_t^u: hour-t regulation-up capacity price [$/kW-h[2]]
p_t^d: hour-t regulation-down capacity price [$/kW-h]
D_t: hour-t building energy demand [kWh]
I_t: hour-t outage state [1 if there is an outage, 0 otherwise]
δ_t^u: hour-$(t-1)$ dispatch-to-contract ratio of regulation up
δ_t^d: hour-$(t-1)$ dispatch-to-contract ratio of regulation down

We also define $S_t = (x_t, p_t^e, p_t^u, p_t^d, D_t, I_t, \delta_t^u, \delta_t^d)$ as a vector of hour-t state variables.

We use the concept of a dispatch-to-contract ratio of regulation up and regulation down to model the relationship between regulation capacity sales and the amount of energy that the battery is obligated to provide as a result of regulation energy being called by the SO [20]. The battery must either charge or discharge energy to meet these obligations or pay the deviation penalty for unserved energy. Kempton and Tomić [20] define the dispatch-to-contract ratio of an AS product as the actual energy dispatched over its contract period divided by the contracted capacity. For example, the hour-t dispatch-to-contract ratio of regulation would be calculated by dividing the total amount of regulation energy dispatched in the hour by the amount of regulation capacity that the SO reserved. Thus, by definition, the terms $\delta_{t+1}^u \cdot k_t^u$ and $\delta_{t+1}^d \cdot k_t^d$ are the amount of regulation-up and -down energy, respectively, that the battery must provide in hour t. Using historical data covering several years, Kempton and Tomić [20] estimate a dispatch-to-contact ratio of 0.08 for regulation in the California ISO market. This value indicates that each kW of regulation capacity sold in the market for one hour will result, on average, in 0.08 kWh of net energy actually being called in real-time. The dispatch-to-contract ratio approach can also be used to model other AS products, such as spinning and non-spinning reserves.

The hour-t ratios are assumed to be unknown when the battery decides how much regulation capacity to sell in hour t.[3] Once the capacity decisions are made, the ratios are then determined and the state of charge of the battery changes or the battery incurs a deviation penalty. We define n_t^u and n_t^d as the amount of regulation-up and -down energy, respectively, that is called by the SO in hour $t-1$ but is unserved. These are auxiliary variables used in our state transition equations and objective function.

[2] A kW-h, which is a unit for one kW of AS capacity provided for one hour, should be distinguished from a kWh, which is a unit of energy.

[3] This is also why we define δ_t^u and δ_t^d as the hour-$(t-1)$ ratios, due to our time definition convention.

3.4 Exogenous Variables

We assume that the variables p_t^e, p_t^u, p_t^d, D_t, I_t, δ_t^u, and δ_t^d evolve randomly and independently of any of the decision variables, but may be dependent on one another. We define $\hat{p}_t^e$, $\hat{p}_t^u$, $\hat{p}_t^d$, $\hat{D}_t$, $\hat{I}_t$, $\hat{\delta}_t^u$, and $\hat{\delta}_t^d$ as exogenous random variables that define the change in the exogenous variables between hour $t-1$ and hour t. These random variables may be dependent on one another. We also define $W_t = (p_t^e, p_t^u, p_t^d, D_t, I_t, \delta_t^u, \delta_t^d)$ as a vector of hour-t exogenous variables. Thus, we can define the vector of hour-t state variables as $S_t = (x_t, W_t)$, where x_t is the endogenous and W_t the exogenous variables. We assume the random process, $\{W_t\}_{t=1}^T$ is Markovian, where T is the optimization horizon. We let $\text{Prob}\{W_{t+1}|W_t\}$ represent the hour-$(t+1)$ transition probabilities. Our solution technique does not require any explicit assumption regarding the dependence of the exogenous random variables.

3.5 State-Transition Function

The state variables, p_t^e, p_t^u, p_t^d, D_t, I_t, δ_t^u, and δ_t^d, evolve randomly according to the following transition equations:

$$\begin{aligned} p_{t+1}^e &= \hat{p}_{t+1}^e + p_t^e, \\ p_{t+1}^u &= \hat{p}_{t+1}^u + p_t^u, \\ p_{t+1}^d &= \hat{p}_{t+1}^d + p_t^d, \\ D_{t+1} &= \hat{D}_{t+1} + D_t, \\ I_{t+1} &= \hat{I}_{t+1} + I_t, \\ \delta_{t+1}^u &= \hat{\delta}_{t+1}^u + \delta_t^u, \end{aligned}$$

and

$$\delta_{t+1}^d = \hat{\delta}_{t+1}^d + \delta_t^d.$$

To define the state-transition function for the storage level of the battery, we first define n_t^u and n_t^d as:

$$n_t^u = \max\left\{0, \delta_t^u k_{t-1}^u - \eta^d (x_{t-1} - \underline{R}) + e_{t-1}^d + e_{t-1}^l - e_{t-1}^c\right\} \tag{1}$$

and

$$n_t^d = \max\left\{0, \delta_t^d k_{t-1}^d - (\overline{R} - x_{t-1})/\eta^c - e_{t-1}^d - e_{t-1}^l + e_{t-1}^c\right\}. \tag{2}$$

These equations define unserved regulation energy as the difference between the amount of regulation called by the SO and the maximum amount of regulation energy that the battery can feasibly provide. Equation (1) defines the maximum amount of regulation-up energy that the battery can provide as the maximum that can be discharged from the battery without violating energy constraints, plus the amount of energy that the battery charges in hour $t-1$ (since the battery can choose not to charge this energy, which provides

more energy in net to the SO). Equation (2) defines the maximum amount of regulation-down energy that the battery can provide as the maximum amount that can be charged plus the amount of energy that the battery discharges in hour $t-1$.

The storage level of the battery evolves according to the function:

$$\begin{aligned} x_{t+1} &= \theta(x_t, W_{t+1}, a_t) \\ &= x_t + \eta^c(e_t^c + \delta_{t+1}^d k_t^d - n_{t+1}^d) - (e_t^d + e_t^l + \delta_{t+1}^u k_t^u - n_{t+1}^u)/\eta^d. \end{aligned}$$

This function defines the storage level at the beginning of hour $t+1$ as the starting storage level at hour t plus net energy charged into the device, less efficiency losses. Energy charged into the battery at hour t equals the sum of energy charged and the amount of regulation down energy that the battery provides. Energy discharged is defined analogously.

3.6 Constraints

Total battery charging and discharging are both constrained to be less than the power capacity of the battery:

$$0 \leq e_t^c + k_t^d \leq \overline{P}^b, \tag{3}$$

and

$$0 \leq e_t^d + e_t^l + k_t^u \leq \overline{P}^b. \tag{4}$$

The storage level of the battery is similarly constrained to be within its upper and lower bounds:

$$\underline{R} \leq x_t \leq \overline{R}. \tag{5}$$

The total net power going into and out of the home is also constrained by the power capacity of the building circuit. Moreover, there can only be net power flows into or out of the house if there is not an outage:

$$-\overline{P}^h(1 - I_t) \leq l_t - e_t^l + e_t^c - e_t^d - k_t^u, \tag{6}$$

and

$$l_t - e_t^l + e_t^c - e_t^d + k_t^d \leq \overline{P}^h(1 - I_t). \tag{7}$$

The battery cannot be charged and arbitrage and regulation sales are not possible if there is an outage:

$$e_t^d, e_t^c, k_t^d, k_t^u = 0, \text{ if } I_t = 1. \tag{8}$$

The building load served in each hour is constrained to be no greater than the building demand:

$$l_t \leq D_t. \tag{9}$$

The decision variables are also constrained to be non-negative:

$$e_t^d, e_t^c, e_t^l, l_t, k_t^u, k_t^d \geq 0. \tag{10}$$

Constraints (6), (7), and (8) together force the building load to be either served by the battery or left unserved in any time period with an outage. We let $\mathcal{A}_s$ denote the set of decision vectors, a, that are feasible in constraints (3) through (10) when the system is in the state s.

3.7 Objective Function

The net profit earned in hour t is given by:

$$C_t(S_t, a_t) = p_t^e(e_t^d - e_t^c) - V^L(D_t - l_t) + p_t^u k_t^u + p_t^d k_t^d \\ + p_{t-1}^e[\delta_t^u k_{t-1}^u - (1 + V^R)n_t^u - \delta_t^d k_{t-1}^d + (1 - V^R)n_t^d].$$

The first term, $p_t^e(e_t^d - e_t^c)$, is the revenue for energy sales. The term, $V^L(D_t - l_t)$, is the penalty associated with unserved building load. The third term, $p_t^u k_t^u + p_t^d k_t^d$, is the payment for regulation capacity sales. The final term, $p_{t-1}^e[\delta_t^u k_{t-1}^u - (1 + V^R)n_t^u - \delta_t^d k_{t-1}^d + (1 - V^R)n_t^d]$, is the revenue for net regulation energy provided, less the penalty for unserved regulation energy.

Let $\mathit{\Pi}$ denote the set of all feasible policies. A policy, $A_t^\pi(S_t)$, is a mapping between an hour-t state, S_t, and a feasible hour-t decision, $a \in \mathcal{A}_s$. For each $\pi \in \mathit{\Pi}$ define the total expected discounted profit from hour t as:

$$G_t^\pi(S_t) = \mathbb{E}\left[\sum_{\tau=t}^{T} \gamma^{\tau-t} C_\tau(S_\tau, A_\tau^\pi(S_\tau)) \middle| S_t\right]. \tag{11}$$

The objective is then to find an optimal policy π^* that satisfies:

$$G_t^{\pi^*}(S_t) = \sup_{\pi \in \mathit{\Pi}} G_t^\pi(S_t),$$

for all $0 \le t \le T$.

4 Solution Technique

Finding exact solutions to our SDP is generally difficult due to the dimension of the problem and the continuous state, decision, and random variables in the model. Thus, we propose solving an approximate version of the SDP in a two-phase manner. In the first phase, the exogenous state and decision variables are discretized, and optimal policies and objective function values for each state are found by solving the discretized SDP (DSDP) using backward induction. In the second phase a mixed-integer program (MIP), in which the value function of the true SDP is approximated as piecewise-linear using the optimal objective function values from the DSDP, is solved to obtain a near-optimal policy. This MIP has a one-hour planning horizon and is solved in a rolling fashion one hour at a time. This provides a feasible policy and can be used to generate a statistical lower bound on the optimal objective function value of the true SDP. We further use a backcasting heuristic to find feasible policies,

which also provides lower bounds, and generate upper bounds using sample average approximation (SAA) and sample path averaging methods. Throughout this discussion we use the notational convention that state and action variables without tildes are from the continuous state and action variable spaces, whereas variables with tildes are from the corresponding discretization.

4.1 Approximation Algorithm

To give the formulation of the DSDP, we let $\tilde{\mathcal{A}}_{\tilde{s}}$ denote the set of decision variables, $\tilde{a}$, that are feasible in constraints (3) through (10) and satisfy the discretization when the system is in state $\tilde{S}$. We define J_t as the set of possible values that $\tilde{S}_t$ can take. We also define Prob$\left\{\tilde{W}_{t+1}|\tilde{W}_t\right\}$ as the transition probabilities of the exogenous state variables in the discretized space. Further details regarding the discretization and the transition probabilities for the discrete exogenous state variables used in our case study are given in Section 5.4.

We can define the following Bellman equation for the DSDP:

$$\begin{aligned}\tilde{F}_t(\tilde{S}_t) &= \max_{\tilde{a}_t \in \tilde{\mathcal{A}}_{\tilde{S}_t}} C_t(\tilde{S}_t, \tilde{a}_t) + \gamma \mathbb{E}\left[\tilde{F}_{t+1}(\tilde{S}_{t+1})|\tilde{S}_t\right], \\ &= \max_{\tilde{a}_t \in \tilde{\mathcal{A}}_{\tilde{S}_t}} C_t(\tilde{S}_t, \tilde{a}_t) \qquad (12) \\ &\quad + \gamma \sum_{j \in J_{t+1}} \text{Prob}\left\{\tilde{W}^j_{t+1}|\tilde{W}_t\right\} \cdot \tilde{F}_{t+1}(\tilde{x}^j_{t+1}, \tilde{W}^j_{t+1}). \qquad (13)\end{aligned}$$

Since we assume that the exogenous random variables are all Markovian, Bellman equation (12) represents a Markov decision process, which can be solved using backward induction. Solving the DSDP gives optimal objective function values, $\tilde{F}^*_t(\tilde{S}_t)$, for each discretized state. These are used in the second phase to construct a MIP, in which a piecewise-linear approximation of the value function of the true SDP is used.

To formulate the MIP, we first define the piecewise-linear approximation of the hour-t value function for a given state, $\tilde{W}^j_t$. This piecewise-linear function has breakpoints at the possible discrete values that $\tilde{x}_t$ can take. If we define M_t as the number of values that $\tilde{x}_t$ can take, then the piecewise-linear approximation for a given state, $\tilde{W}^j_t$, is given by:

$$F^{pl}_t(x_t, \tilde{W}^j_t) = \sum_{m=1}^{M_t-1} \frac{\tilde{F}^*_t(\tilde{x}^{m+1}_t, \tilde{W}^j_t) - \tilde{F}^*_t(\tilde{x}^m_t, \tilde{W}^j_t)}{\tilde{x}^{m+1}_t - \tilde{x}^m_t} \cdot q^{m,j}_t,$$

where:

$$\sum_{m=1}^{M_t-1} q^{m,j}_t = x_t,$$

$$0 \le q^{m,j}_t \le y^{m,j}_t \cdot (\tilde{x}^{m+1}_t - \tilde{x}^m_t), \forall\ m = 1, \ldots, M_t - 1,$$

and

$$y_t^{m,j} \leq y_t^{m+1,j} \in \{0,1\}, \forall\ m = 1, \ldots, M_t - 1.$$

The auxiliary variable, $y_t^{m,j}$, indicates which piece of the approximation x_t lies in, and $q_t^{m,j}$ indicates how much storage capacity in each piece x_t covers. For ease of notation, we can also define the piecewise linear approximation as:

$$F_t^{pl}(x_t, \tilde{W}_t^j) = \sum_{m=1}^{M_t - 1} \tilde{\sigma}_t^m(\tilde{W}_t^j) \cdot q_t^{m,j},$$

where:

$$\tilde{\sigma}_t^m(\tilde{W}_t^j) = \frac{\tilde{F}_t^*(\tilde{x}_t^{m+1}, \tilde{W}_t^j) - \tilde{F}_t^*(\tilde{x}_t^m, \tilde{W}_t^j)}{\tilde{x}_t^{m+1} - \tilde{x}_t^m},$$

are the slopes of the piecewise-linear approximation. The formulation of the hour-t MIP is then given by:

$$\max_{a_t \in \mathcal{A}_{S_t}} F_t^{MIP}(S_t, a_t) = \max_{a_t \in \mathcal{A}_{S_t}} C_t(S_t, a_t) + \gamma \sum_{j \in J_{t+1}} \text{Prob}\left\{\tilde{W}_{t+1}^j | W_t\right\} \cdot F_{t+1}^{pl}(\theta(x_t, \tilde{W}_{t+1}^j, a_t), \tilde{W}_{t+1}^j). \tag{14}$$

This MIP uses a piecewise linear approximation of the true cost-to-go function of the SDP, which is constructed from the optimal value of the DSDP. The expected cost-to-go, $\mathbb{E}\left[F_{t+1}(S_{t+1})|S_t\right]$, is computed numerically by averaging over the set of discrete values that $\tilde{W}_{t+1}$ is assumed to take in the DSDP. Given an exogenous hour-t state, W_t, from the continuous space, we compute conditional transition probabilities, $\text{Prob}\left\{\tilde{W}_{t+1}^j | W_t\right\}$, to the discrete states, $\tilde{W}_{t+1}^j$, used in the DSDP. This is generally done by discretizing the conditional distribution $\text{Prob}\{W_{t+1}|W_t\}$, and we describe in detail how this is in our case study in Section 5.4. The term, $\theta(x_t, \tilde{W}_{t+1}^j, a_t)$, represents the hour-$(t+1)$ storage level resulting from the chosen action, a_t, and the realized exogenous state, $\tilde{W}_{t+1}^j$, We do not restrict x_{t+1} to take on a discretized value, since we interpolate on the storage variable. This also allows the action variables to be chosen from the continuous space. Alternatively, one could also attempt interpolation on the exogenous state variables. Our results in Section 6.5 suggest, however, that interpolation on the storage level variable only is sufficient, insomuch as this method provides policies with relatively small optimality gaps. Note that this MIP only determines hour-t actions—our algorithm uses a forward-rolling optimization method to find a policy over the full planning horizon. The detailed formulation of the MIP, including all of the constraints, is given in Appendix A.

An optimal solution of the MIP is feasible but not optimal in the true SDP. This is because the only factors that are unknown when making hour-t decisions and affect the system state at hour $t+1$ are the dispatch-to contract-ratios. However, since regulation calls can go unserved, with a penalty, unserved regulation energy can be defined accordingly to ensure feasibility of

the solution. The pseudocode in Algorithm 1 summarizes the approximation algorithm that we use to find near-optimal solutions to our SDP. Steps 1 and 2 represent the first phase of the algorithm, in which the SDP is discretized and the resulting DSDP is solved exactly using backward induction. The second phase of the algorithm works by iterating through the hours of the optimization horizon. In each hour, the exogenous variables, W_t, are first observed (step 5). Then the amount of unserved regulation energy in hour $t-1$ is updated, based on the actual hour-$(t-1)$ dispatch-to-contract ratio (step 7) and the resulting energy level of the battery is determined (step 8). Finally, the MIP is solved to determine the hour-t actions, $\hat{a}_t$, (step 10) and the hour-t profit contribution is calculated (step 11).

Algorithm 1 Approximation Algorithm Pseudocode

1: Initialize: Discretize state and action variables x_t, p_t^e, p_t^u, p_t^d, D_t, I_t, δ_t^u, δ_t^d, e_t^d, e_t^c, e_t^l, l_t, k_t^u, and k_t^d
2: Solve DSDP exactly using backward induction and obtain optimal objective function values, $\tilde{F}_t^*(\tilde{S}_t)$, for each discretized state
3: Fix x_1 equal to lowest discretized value of x_t {assume battery starts empty}
4: **for** $t = 1$ to T **do**
5: Sample W_t from continuous random distribution
6: **if** $t > 1$ **then**
7: Let $n_t^u \leftarrow \max\left\{0, \delta_t^u \hat{k}_{t-1}^u - \eta^d(x_{t-1} - \underline{R}) + \hat{e}_{t-1}^d + \hat{e}_{t-1}^l - \hat{e}_{t-1}^c\right\}$ and $n_t^d \leftarrow \max\left\{0, \delta_t^d \hat{k}_{t-1}^d - (\overline{R} - x_{t-1})/\eta^c - \hat{e}_{t-1}^d - \hat{e}_{t-1}^l + \hat{e}_{t-1}^c\right\}$ {update hour-$(t-1)$ unserved regulation}
8: Let $x_t \leftarrow x_{t-1} + \eta^c(\hat{e}_{t-1}^c + \delta_t^d \hat{k}_{t-1}^d - n_t^d) - (\hat{e}_{t-1}^d + \hat{e}_{t-1}^l + \delta_t^u \hat{k}_{t-1}^u - n_t^u)/\eta^d$ {update hour-t storage level}
9: **end if**
10: Let $\hat{a}_t \leftarrow \arg\max_{a_t \in \mathcal{A}_{S_t}} F_t^{MIP}(S_t, a_t)$ {solve the MIP}
11: Let $C_t \leftarrow p_t^e(\hat{e}_t^d - \hat{e}_t^c) - V^L(D_t - \hat{l}_t) + p_t^u \hat{k}_t^u + p_t^d \hat{k}_t^d + p_{t-1}^e[\delta_t^u \hat{k}_{t-1}^u - (1+V^R)\hat{n}_t^u - \delta_t^d \hat{k}_{t-1}^d + (1-V^R)\hat{n}_t^d]$ {calculate hour-t profit contribution}
12: **end for**

4.2 Statistical Bounds

To evaluate the quality of the policies found by the approximation algorithm, we compute two statistical lower and upper bounds on the optimized value of the SDP. We say that an estimator, $\hat{V}_N$, is a valid statistical upper bound of the true optimal value, V^*, if $V^* \leq \mathbb{E}\left[\hat{V}_N\right]$, and similarly, $\hat{V}_N$ is a valid statistical lower bound if $\mathbb{E}\left[\hat{V}_N\right] \leq V^*$, where N is the number of replications of the random variables underlying the calculation of $\hat{V}_N$. The statistical bound, $\hat{V}_N$, is consistent if $\hat{V}_N \to V^*$ as $N \to \infty$.

We use implementable and feasible policies to generate lower bounds on the optimized objective value of the SDP. A policy is implementable and feasible if it satisfies the constraints of the problem and is nonanticipative. The

resulting objective function value from using such a policy provides a statistical lower bound on the optimal value of the true SDP [29]. One of the lower bounds that we compute is found by randomly generating sample paths, ω, of the exogenous random variables, W, and using the approximation algorithm, outlined in Section 4.1, to derive a feasible policy. Because the approximation algorithm assumes that the hour $s > t$ exogenous random variables are unknown when hour-t decisions are made, the policy is nonanticipative. Let $G^{IF}(\omega^i)$ denote the profit of the battery over the T-hour horizon given the sample path, ω^i, of exogenous state variables. This is computed as the sum of the profit contributions calculated in step 11 of Algorithm 1 when using an implementable and feasible policy derived using the algorithm (this also gives rise to the superscript, IF, indicating that the bound is generated using implementable and feasible policies). Note that step 11 of Algorithm 1 computes the profit contributions using the 'true' values of the exogenous state variables, drawn from the continuous space. Thus, $G^{IF}(\omega^i)$ is a lower bound on the value of the optimal objective function value of SDP over the T-hour horizon, given ω^i. The statistical lower bound using $\mathcal{I}$ sample paths is given by:

$$B_L^{IF} = \frac{1}{\mathcal{I}} \sum_{i=1}^{\mathcal{I}} G^{IF}(\omega^i).$$

We generate 100 of these statistical lower bounds, using a random sample of 1,000 i.i.d (independent and identically distributed) ω^i's for each. By generating 100 replications of the lower bound, we compute a sample standard error, providing confidence intervals.

Our other statistical lower bound is also generated from implementable and feasible policies, which are derived using a backcasting heuristic [35,34]. The pseudocode in Algorithm 2 outlines the heuristic, which relies on the fact that the exogenous random variables tend to exhibit diurnal patterns (we explicitly model these types of patterns in our case study, as discussed in Section 5). The heuristic determines decisions in a forward-rolling fashion. In each hour, the exogenous state variables, W_t, are observed (step 2) and the starting storage level of the battery is updated based on the amount of regulation energy called in the previous hour (steps 4 and 5). The exogenous random variables in hours $t+1$ through $t+23$ are assumed to equal the actual values from the previous day (step 7), and a deterministic linear program with a 24-hour optimization horizon is solved (step 8). Appendix B gives the detailed formulation of this linear program. The hour-t action variables from the linear program solution are used (step 9) and the hour-t profit contribution is computed based on these actions (step 10).

As with our approximation algorithm, the backcasting heuristic generates implementable and feasible policies. This is because hour-t decisions are determined based solely on historical information that is available at that time. The policy is feasible because we compute the unserved regulation energy based on the actual dispatch-to-contract ratio in each hour. Let $G^{BC}(\omega^i)$ denote the profit of the battery over the T-hour horizon given the sample path, ω^i (the

Algorithm 2 Backcasting Heuristic Pseudocode

1: Let $x_1 \leftarrow \underline{R}$ {assume battery starts empty}
2: **for** $t = 1$ to T **do**
3: Sample W_t from continuous random distribution
4: **if** $t > 1$ **then**
5: Let $n_t^u \leftarrow \max\left\{0, \delta_t^u \hat{k}_{t-1}^u - \eta^d(x_{t-1} - \underline{R}) + \hat{e}_{t-1}^d + \hat{e}_{t-1}^l - \hat{e}_{t-1}^c\right\}$ and $n_t^d \leftarrow \max\left\{0, \delta_t^d \hat{k}_{t-1}^d - (\overline{R} - x_{t-1})/\eta^c - \hat{e}_{t-1}^d - \hat{e}_{t-1}^l + \hat{e}_{t-1}^c\right\}$ {update hour-$(t-1)$ unserved regulation}
6: Let $x_t \leftarrow x_{t-1} + \eta^c(\hat{e}_{t-1}^c + \delta_t^d \hat{k}_{t-1}^d - n_t^d) - (\hat{e}_{t-1}^d + \hat{e}_{t-1}^l + \delta_t^u \hat{k}_{t-1}^u - n_t^u)/\eta^d$ {update hour-t storage level}
7: **end if**
8: Let $W_\tau \leftarrow W_{\tau-24}\ \forall\ \tau = t+1, \ldots, t+23$ {assume previous day's exogenous random variables repeat themselves}
9: Solve the 24-hour deterministic linear program given in B, let $\{\dot{a}_\tau\}_{\tau=t}^{t+23}$ denote optimized values of the decision variables
10: Let $\hat{a}_t \leftarrow \dot{a}_t$ {save hour-t decision variables from the deterministic linear program solution only}
11: Let $C_t \leftarrow p_t^e(\hat{e}_t^d - \hat{e}_t^c) - V^L(D_t - \hat{l}_t) + p_t^u \hat{k}_t^u + p_t^d \hat{k}_t^d + p_{t-1}^e[\delta_t^u \hat{k}_{t-1}^u - (1+V^R)\hat{n}_t^u - \delta_t^d \hat{k}_{t-1}^d + (1-V^R)\hat{n}_t^d]$ {calculate hour-t profit contribution}
12: **end for**

superscript, BC, indicates that this bound is generated using the backcasting technique). This is given by the sum of profits over the T-hour horizon using the policy derived by using Algorithm 2, as computed in step 11 of the algorithm. The statistical lower bound generated by applying the backcasting heuristic to $\mathcal{I}$ sample paths is given by:

$$B_L^{BC} = \frac{1}{\mathcal{I}} \sum_{i=1}^{\mathcal{I}} G^{BC}(\omega^i).$$

As before, we generate 100 of these bounds using a random sample of 1,000 ω^i's for each.

We generate one set of upper bounds using a two-stage SAA technique. This is done by solving a stochastic program, the formulation of which is given in Appendix C, in which the expected profit of the battery is maximized with uncertain future exogenous random variables. Specifically, the stochastic program has a two-stage scenario tree, in which the first stage is hour 1 and the second stage is hours 2 through T. Thus, the scenario tree assumes that hour-1 decisions are made without knowledge of the exogenous random variables in hours 2 through T, while subsequent decisions are made with full knowledge of future states. Let B_U^{SAA} denote the optimal expected objective function value of the problem in stage one (the superscript, SAA, indicates that this bound is generated using an SAA). B_U^{SAA} provides a valid statistical upper bound on the optimized objective function value of the true SDP with a T-hour planning horizon [29]. The scenario tree used has 20 sample paths and is randomly generated by sampling from the exogenous random variable distributions.[4]

[4] The scenario tree only has 20 sample paths because the resulting two-stage model is computationally intractable with a larger tree.

We generate and solve 1,000 of these stochastic programs in order to compute a standard error for the upper bound.

Our other set of upper bounds is generated using a sample path averaging technique. This bound is computed by randomly generating sample paths of the exogenous random variables. For each sample path, ω^i, of exogenous state variables we solve a deterministic linear program in which the full sequence of random variable realizations is known in hour 1. The linear program has the same structure as that used for the backcasting heuristic, and the formulation is given in Appendix B. Let $G^{DET}(\omega^i)$ denote the profit of the battery over the T-hour horizon with the sample path, ω^i (the superscript, DET, indicates that this bound is generated using deterministic problems). The statistical upper bound generated by applying the averaging technique to $\mathcal{I}$ sample paths is then given by:

$$B_U^{DET} = \frac{1}{\mathcal{I}} \sum_{i=1}^{\mathcal{I}} G^{DET}(\omega^i).$$

We generate 100 of these bounds using a random sample of 1,000 ω^i's for each.

SAA gives a consistent statistical upper bound when the random variables are mutually and serially independent [29]. Otherwise, applying a conditional sampling scheme yields consistency with dependent random variables. In our case study, the random variables, I_t, are serially dependent. We do not apply a conditional sampling scheme, however, since even two conditional samples at each stage yields 2^{167} sample paths, making the problem intractable. Hence, our SAA technique does not guarantee consistent upper bounds.

5 Case Study

To provide a numerical example of our model we study a case in which a single battery is installed in a home in the state of Ohio. We assume that the home is in the PJM Interconnection system, and use PJM market rules and data to set the parameter values and the distribution of the random variables. In doing so, we use market and system data from the summer of 2009 to study battery use and values over a 'typical' summer week (a 168-hour horizon).

5.1 Battery Characteristics

We assume that the battery has an energy storage capacity of $\overline{R} = 11.2$ kWh and a minimum energy level of $\underline{R} = 3$ kWh. The battery is also assumed to have a charging and discharging capacity of $\overline{P}^b = 7.2$ kW, which implies that it is connected to a 240 V/30 A appliance circuit. We assume 10% efficiency losses when charging and discharging the battery, meaning $\eta^c = \eta^d = 0.9$. This further implies that the battery has a roundtrip efficiency of $\eta^c \cdot \eta^d = 0.81$. The building is assumed to have a power capacity of $\overline{P}^h = 10$ kW, which is typical of homes in the region that we model.

5.2 Market Assumptions

The PJM market treats regulation up and down as a single capacity product with a single price, and we use this convention in our case study. In terms of our model, this means that the state variables, p_t^u and p_t^d, are fixed equal to each other in every hour (we define the value of these state variables as $p_t^u = p_t^d = p_t^r/2$, where p_t^r is the regulation capacity price in hour t, so as not to double-count capacity payments). The decision variables, k_t^u and k_t^d, are also constrained to equal each other in every hour, since the battery cannot differentiate between sales of regulation-up and -down capacity. We let $k_t = k_t^u = k_t^d$ denote the amount of regulation capacity sold in hour t. The dispatch-to-contract ratios, δ_t^u and δ_t^d, are not fixed equal to one another, because the system may need more regulation energy in one direction than the other.

We assume that an unserved regulation deployment by the battery incurs a penalty of $V^R = 0.15$, based on PJM tariffs. This means that if the battery cannot supply regulation-up energy, it must purchase replacement energy from the market at 115% of the real-time energy price. Likewise, if it cannot provide regulation-down energy, then it must sell its excess energy to the market at 85% of the real-time energy price. We use a penalty cost of $V^L = \$3.72/\text{kWh}$ on any unserved building load [19]. Because we only model a single week, we assume a discount factor of $\gamma = 1$.

5.3 Exogenous Random Variables

Although our model does not require any specific correlation structure among the exogenous random variables, $\hat{p}_t^e$, $\hat{p}_t^u$, $\hat{p}_t^d$, $\hat{D}_t$, $\hat{I}_t$, $\hat{\delta}_t^u$, and $\hat{\delta}_t^d$, are all assumed to be mutually independent in our case study. We further assume that the price, load, and dispatch-to-contact ratio variables are serially independent through time. These assumptions are based on empirical analyses of historical PJM and Ohio power system data.

Historical energy and regulation capacity prices have little correlation—*e.g.*, during 2009 these prices had a correlation of about -0.15. This is because high energy prices signal less generating capacity being available and higher-cost generation having to be used to serve the load, whereas high regulation prices signal a lack of fast-responding generation. Indeed, energy prices tend to peak during the day when electricity demand is the highest, whereas regulation prices can peak overnight when loads are low and only baseload generators with slow ramping rates are online. These differences in the diurnal price patterns also explain the slightly negative correlation between energy and regulation prices. Tests of historical price data from the California ISO and Singapore electricity markets show that energy price patterns best fit a log-normal distribution [21,30]. Our own examination of 2009 PJM energy prices show similar results, and as such we assume that the energy prices have a log-normal distribution. Although the distribution of regulation prices has not been examined, our analysis of 2009 PJM data shows that these prices

also fit a log-normal distribution, which we assume. In order to capture diurnal energy and regulation price patterns, we allow for different location and scale parameters in the log-normal distributions for each of the 24 hours of the day. We fit these parameters using least-square estimation based on price data from the PJM market in the summer of 2009.

Seppala [28] examines the statistical properties of the electricity demand of residential homes. He compares several parametrized distribution functions and finds that a log-normal distribution provides the best fit. Thus, we assume that the building demand has a log-normal distribution and allow the location and scale parameters to vary in each of the 24 hours of the day. We fit these parameters using least-square estimation based on historical residential load data for the summer of 2009 provided by AEP for a set of its customers in Ohio. The loads correspond to a home that is approximately 200 m^2 (2200 ft^2) in size. Although there is a relationship between energy prices and loads, we are modeling a single building which has only a marginal effect on the system. Moreover, since we allow the distribution of the hourly prices and loads to vary, this captures any coincidence in energy prices and building demand.

The distributions of the dispatch-to-contract ratios for regulation up and down are estimated using least-square estimation based on historical PJM data from the summer of 2009. These data specify the amount of regulation capacity reserved in each hour and the amount of regulation energy deployed in real-time. The data do not show any diurnal patterns in the ratios. As such, we assume that the distributions are time-invariant. Hypothesis testing shows that a Gaussian distribution best fits the historical data, which we use in our case study.

A number of approaches are used to model power system reliability, with Markov-based models being the most common. The mechanics of component failure and repair suggest that power system failure follows a Markov process with exponentially distributed time between failures [8]. Semi-Markov models are also used, however [2,4]. We model system outages using a two-state Markov chain in which the system can either be in an outage or non-outage state. Transitions between these states depend solely on the present state of the system, and the transition probabilities are time-invariant. We use system reliability data reported by FirstEnergy in 2009, which had an average of 1.24 outages per customer during the year which lasted an average of 2 hours, to estimate the transition probabilities. Based on these values, we assume a probability of 0.000142 that the system has an outage at hour $t+1$ if it is in a non-outage state at hour t, and a probability of 0.5 that it recovers from an outage in each hour.

5.4 Model Discretization and Algorithm Implementation

Table 1 summarizes the discretization of the state variables used in our approximation algorithm. The underlying distributions of the $\hat{p}_t^e$, $\hat{\delta}_t^u$, and $\hat{\delta}_t^d$ variables are discretized into five possible outcomes using bracket medians.

The distributions of the $\hat{p}_t^u$, $\hat{p}_t^d$ variables are similarly discretized into four possible outcomes and the distribution of the $\hat{D}_t$ variables into three possible outcomes. The outage variable, I_t, is assumed to only take on two values, thus no further discretization of this variable is done. The storage level state variable, x_t, is discretized into 21 possible values with equal interval widths between $\underline{R}$ and $\overline{R}$. Because we assume that the regulation capacity prices, p_t^u and p_t^d, equal each other in each hour, this eliminates one dimension of the state space giving 63,000 possible states in each hour.

Table 1 Discretization of state variables in approximation algorithm.

Variable	Type	Distribution	Number of Values
$\tilde{p}_t^e$	Exogenous	Log-normal	5
$\tilde{p}_t^u = \tilde{p}_t^d$	Exogenous	Log-normal	4
$\tilde{D}_t$	Exogenous	Log-normal	3
I_t	Exogenous	Bernoulli	2
$\tilde{\delta}_t^u$	Exogenous	Gaussian	5
$\tilde{\delta}_t^d$	Exogenous	Gaussian	5
$\tilde{x}_t$	Endogenous	n/a	21

Since we use bracket medians to discretize the price, building demand, and dispatch-to-contract ratio variables, each discretized value is equally likely. Moreover, since we assume that these variables and the outage state are mutually independent, we have the following state-transition probabilities:

$$\begin{aligned}\text{Prob}\left\{\tilde{W}_{t+1}|W_t\right\} &= \text{Prob}\left\{\tilde{p}_{t+1}^e\right\} \cdot \text{Prob}\left\{\tilde{p}_{t+1}^u (= \tilde{p}_{t+1}^d)\right\} \cdot \text{Prob}\left\{\tilde{D}_{t+1}\right\} \\ &\quad \cdot \text{Prob}\left\{I_{t+1}|I_t\right\} \cdot \text{Prob}\left\{\tilde{\delta}_{t+1}^u\right\} \cdot \text{Prob}\left\{\tilde{\delta}_{t+1}^d\right\} \\ &= \frac{1}{5} \cdot \frac{1}{4} \cdot \frac{1}{3} \cdot \text{Prob}\left\{I_{t+1}|I_t\right\} \cdot \frac{1}{5} \cdot \frac{1}{5} \\ &= \frac{1}{1500} \cdot \text{Prob}\left\{I_{t+1}|I_t\right\},\end{aligned}$$

where the term, $\text{Prob}\left\{\tilde{p}_{t+1}^u (= \tilde{p}_{t+1}^d)\right\}$, highlights the fact that we restrict the regulation prices to equal each other in each hour and $\text{Prob}\left\{I_{t+1}|I_t\right\}$ depends on the hour-t outage state.

The assumption that $\tilde{x}_t$ takes on only 21 possible values implies that the charging and discharging variables, $\tilde{e}_t^d$, $\tilde{e}_t^c$, and $\tilde{e}_t^l$, can take at most 21 values, corresponding to possible state transitions between $\tilde{x}_t$ and $\tilde{x}_{t+1}$. We discretize the regulation capacity variables, $\tilde{k}_t$, using 1 kW interval widths, which implies that these variables can take on at most 8 values. Table 2 summarizes the number of possible values that the state variables can take in our discretization. Again, because we assume that the action variables $\tilde{k}_t^u$ and $\tilde{k}_t^d$ must equal each other in each hour, this eliminates one dimension of the action variable.

Table 2 Discretization of action variables in approximation algorithm.

Variable	Number of Discretized Values
$\tilde{e}_t^d$	21
$\tilde{e}_t^c$	21
$\tilde{e}_t^l$	21
$\tilde{l}_t$	21
$\tilde{k}_t = \tilde{k}_t^u = \tilde{k}_t^d$	8

When solving the DSDP, we go through the action space and enumerate all feasible combinations of action variables. However, the actual number of values that these variables can take is significantly less than the product of the values in Table 2, due to the power capacity constraints. Moreover, as it is suboptimal to charge and discharge energy simultaneously for non-regulation sales (due to the efficiency losses), we can eliminate action vectors in which both $\tilde{e}_t^c$ and $\tilde{e}_t^d$ are simultaneously positive from consideration. The starting storage level, $\tilde{x}_t$, and the outage state, I_t, also limit the number of possible state transitions. Algorithm 3 summarizes the steps used for such action enumeration, which reduces the total number of feasible actions in each hour to 170 (18 when there is an outage and 152 otherwise). In the algorithm, we define Δ as the width of the intervals between the discrete values which the $\tilde{e}_t^d$, $\tilde{e}_t^c$, and $\tilde{e}_t^l$ variables can take. In our case study, $\Delta = 0.41$ kW. Note that due to the high penalty on unserved building loads, the value of $\tilde{l}_t$ can be determined by constraints (6), (7), and (9) and the values of I_t, $\tilde{e}_t^c$, $\tilde{e}_t^d$, $\tilde{e}_t^l$, and $\tilde{k}_t$. Specifically, if $I_t = 1$, then $\tilde{l}_t = \min\{\tilde{D}_t, \tilde{e}_t^l\}$. Otherwise, if $I_t = 0$ and $\tilde{D}_t \leq \overline{P}^h$, then $\tilde{l}_t = \tilde{D}_t$. These two conditions are true due to the high penalty on unserved building loads. Finally, if $I_t = 0$ and $\tilde{D}_t > \overline{P}^h$ then we let:

$$\tilde{l}_t = \overline{P}^h + \tilde{e}_t^l.$$

6 Case Study Results

We study the interactions between and values of different storage applications by examining four cases, which allow the battery to be used for different storage applications. We examine (i) arbitrage-only; (ii) arbitrage and backup-energy; (iii) arbitrage, backup-energy, and regulation; and (iv) arbitrage, backup-energy, regulation, and distribution-relief cases. Since storage use depends on the realization of the exogenous random variables, we use a common randomly generated sample path (which is drawn from the continuous distributions) in our comparison between the four different cases.

Algorithm 3 State Reduction Procedure Pseudocode

for $t = 1$ to T **do**
  if $I_t = 0$ **then**
    for $\tilde{k}_t = 0$ to 7 **do**
      Let $\alpha_t \leftarrow \overline{R} - \tilde{k}_t$ {remaining power capacity for services other than regulation}
      for $\tilde{e}_t^c = \Delta$ to α_t in increments of Δ **do**
        Let $\tilde{e}_t^d, \tilde{e}_t^l \leftarrow 0$ {suboptimal to charge and discharge simultaneously}
      end for
      for $\tilde{e}_t^d = \Delta$ to α_t in increments of Δ **do**
        for $\tilde{e}_t^l = \Delta$ to $\alpha_t - \tilde{e}_t^d$ in increments of Δ **do**
          Let $\tilde{e}_t^c \leftarrow 0$ {suboptimal to charge and discharge simultaneously}
        end for
      end for
    end for
  else
    for $\tilde{e}_t^l = 0$ to $\overline{P}^b$ in increments of Δ **do**
      Let $\tilde{e}_t^d, \tilde{e}_t^c, \tilde{k}_t \leftarrow 0$ {no grid services during outage}
    end for
  end if
end for

6.1 Arbitrage-Only

We first examine a case in which the battery is used solely for energy arbitrage. In this case we assume that the system does not experience outages (*i.e.*, we fix $I_t = 0$ with probability 1 for all t), that the battery cannot provide regulation (*i.e.*, we fix $k_t^u, k_t^d = 0 \; \forall \; t$), and that the distribution system is not overloaded (*i.e.*, $D_t \leq \overline{P}^h$ with probability 1 for all t). Figure 1 shows hourly energy prices and net energy sales by the battery in the arbitrage-only case on a single day. Because the battery is used solely for arbitrage, the value generated by the battery is relatively low—the actual realized value of the battery over the week amounts to about $1.85.

6.2 Backup Energy

In the backup-energy case we relax the restriction that I_t be fixed equal to zero in every hour and allow the system to experience outages according to the probabilities given in Section 5.3. We still assume in this case that the battery cannot provide regulation and that the distribution system is not overloaded by the building load. Figure 2 shows the starting storage level of the battery in the arbitrage-only and arbitrage and backup cases. The results are for the same day shown in Figure 1. Comparing the two cases shows that the storage level of the battery is kept higher in the backup-energy case, in order to provide a safety stock of energy in case of an outage. This is done at the expense of arbitrage profits since the battery is charged using higher-priced energy in hour 3.

These battery usage patterns persist over the week—on average about 6.8 kWh of energy is kept in storage in each hour in the backup-energy case as

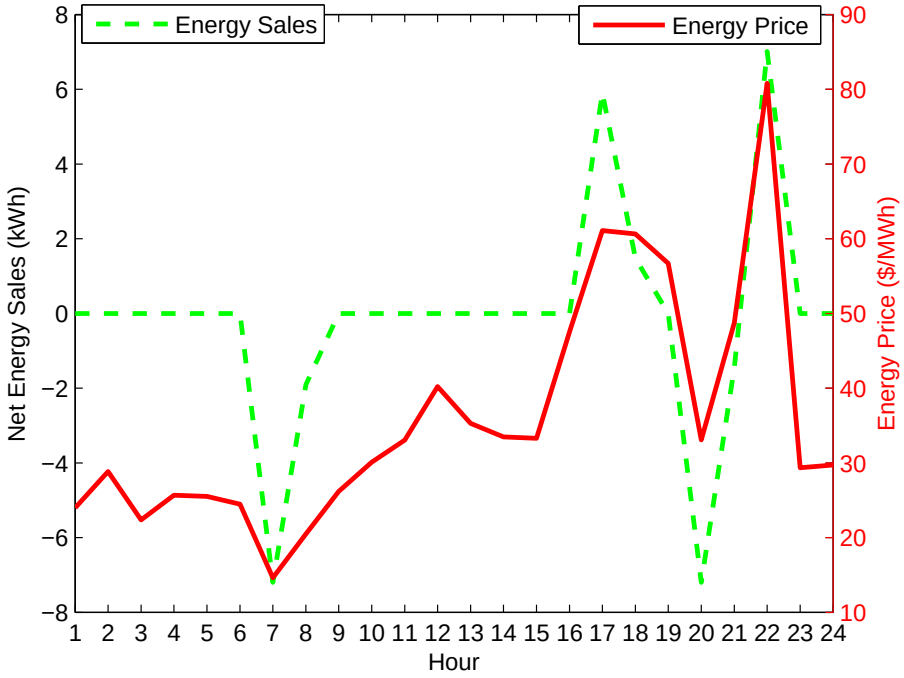

Fig. 1 Net energy discharged from battery in arbitrage-only case.

opposed to only 6.6 kWh in the arbitrage-only case. This higher storage level in the backup-energy case translates into a very slight reduction in arbitrage profits of about a cent, since the battery does less arbitrage. The energy in the battery provides about $0.20 of expected benefit in averting the cost of a system outage (there are no outages during the week that we simulate, thus we report the expected benefit based on the expected cost of lost load in each hour). If the optimized policy from the arbitrage-only case is used in a system in which outages can occur, the energy in storage provides a benefit of about $0.18 in averting expected lost load. Thus optimized battery use in the backup energy case forgoes less than a cent of arbitrage profit to gain about two cents of expected backup energy value.

6.3 Regulation Services

In this case we allow the battery to provide regulation, by relaxing the constraint that $k_t^u, k_t^d = 0 \ \forall \ t$, as well as arbitrage and backup-energy services. When the battery can provide regulation it almost exclusively provides this service and very little arbitrage. This is because regulation is primarily a capacity service with little net charging of energy, meaning that providing this service incurs very little cost and comparably high revenue. Although the dispatch-to-contract ratio in a particular hour can be high—we find cases of

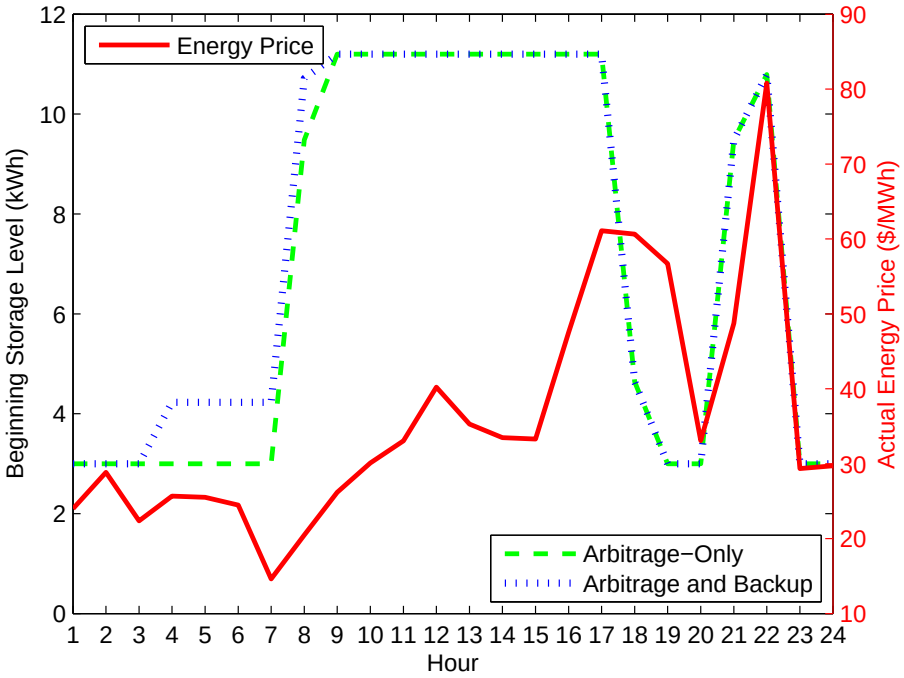

Fig. 2 Starting storage level of battery in arbitrage-only and arbitrage and backup-energy cases.

up to 0.35 in the historical PJM data—the regulation-up and -down signals tend to cancel out in the long run. Our simulation has high ratios of up to 0.30, but the average ratio over the week is much lower at -0.06, which is consistent with the historical PJM data. Thus, on average, providing regulation results in small net charging of the battery. This use of the battery reduces arbitrage profits even further compared to the other two cases—the battery earns $1.25 over the course of the week when regulation services are allowed—but the regulation profits of $23.90 more than compensate for this.

6.4 Distribution Relief

In this case we assume that an added plug-in hybrid electric vehicle (PHEV) charging load can overload the $\overline{P}^h = 10$ kW power constraint of the electric circuit connected to the home. We assume that the on-board PHEV battery has an energy capacity of 14 kWh and is charged using a 7.2 kW appliance circuit when parked at home. The PHEV driving and charging patterns are assumed to be fixed and deterministic, and are based on actual vehicle use data taken from an empirical fleet study [22]. The amount of battery energy used when the PHEV is driven, which determines the overnight charging load, is estimated using a dynamic vehicle simulator [38].

Figure 3 shows the hourly load profile of the home with and without the PHEV. The loads are shown for two days on which the total load of the home exceeds the 10 kW capacity with the addition of the PHEV. Figure 3 also shows the starting storage level of the battery in each hour with and without the PHEV. When the PHEV is added, extra energy is stored in the battery to relieve the building power constraint, which is violated in hours 22 and 44. This use of the battery allows the home to avoid $18.81 in costs associated with loads that would otherwise not be served due to the distribution constraint during the week. As with the backup energy case, this reduces the market value of the battery due to reduced arbitrage and regulation profits. For instance, to keep a high SOC the battery buys an additional 2.7 kWh of energy during the two days shown when the PHEV load is added. Moreover, the battery provides 2 kW-h less regulation capacity in hours 22 and 44, since the higher net load of the home may prevent it from providing regulation down in real-time.

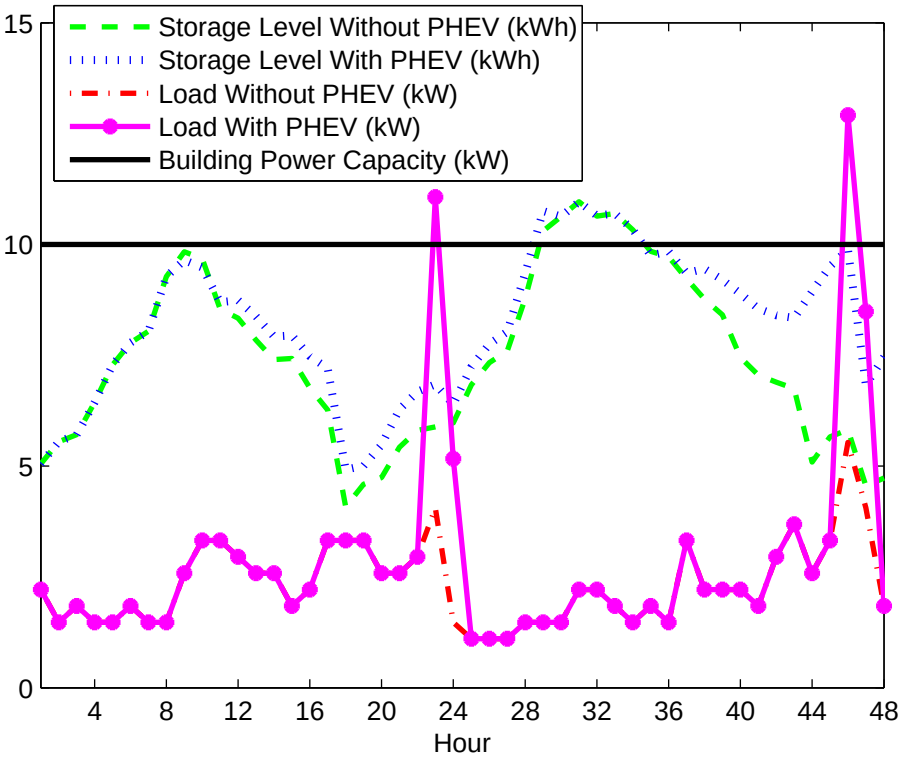

Fig. 3 Total load of home and starting storage level of battery with and without PEV.

6.5 Solution Quality

Table 3 summarizes the mean upper and lower bounds on the optimal value of the SDP for the four storage application cases that we consider. It also provides standard errors for the bounds and an average optimality gap bound. The optimality gap bound is defined as the difference between the mean upper and lower bounds, as a percent of the lower bound. Since the expected averted cost savings from providing backup energy and distribution relief is not included in the objective function (rather, these costs are subtracted if incurred), the objective function values decrease between the arbitrage-only and backup energy and the regulation and distribution relief cases. The optimality gaps are relatively small, showing that our approximation algorithm provides relatively good near-optimal policies. Moreover, the gaps tend to be smaller in the cases in which regulation service is provided. This is because the high profits from regulation imply that an optimal policy is relatively insensitive to accurately forecasting future prices. The gaps are higher in the arbitrage-only and backup energy cases since arbitrage profits are more sensitive to accurately estimating energy price patterns. On the other hand, the arbitrage-only and backup energy cases are less complex than the regulation models, thus a finer discretization could be used which could provide better policies with little incremental computational cost.

Table 3 Upper and lower bounds on optimized value of week-long SDP.

	Arbitrage-Only	Backup Energy	Regulation	Distribution Relief
B_L^{IF}				
Mean	1.77	1.76	25.14	24.47
SE	0.0034	0.0034	0.0065	0.0077
B_L^{BC}				
Mean	0.86	0.86	24.87	23.33
SE	0.0864	0.0864	0.2876	0.6419
B_U^{SAA}				
Mean	1.82	1.82	25.42	24.73
SE	0.128	0.128	0.342	0.407
B_U^{DET}				
Mean	1.82	1.82	25.42	24.73
SE	0.0021	0.0021	0.0075	0.0069
Gap [%]	2.6	3.1	1.1	1.0

The results also show that the backcasting heuristic performs worse than our approximation algorithm. In cases with regulation, the optimality gaps are reasonable, ranging between 2% and 6%. In the other cases, however, the optimality gaps are greater than 100%. This suggests that the backcasting

heuristic is considerably less reliable in cases involving energy arbitrage only. Nevertheless, the backcasting heuristic is significantly less computationally expensive than our approximation algorithm, since the method requires only solving a deterministic linear program. Thus the method may be valuable, even in cases involving energy arbitrage only, due to the simplicity of implementing it.

7 Conclusions

This paper introduces an SDP model that co-optimizes the use of a storage device that is put to multiple uses, such as energy arbitrage, AS, backup energy, and distribution relief. The model is applied to a case study in which a battery is installed in a residential home. Although our example assumes a small battery used as distributed storage, the model we develop can be applied to other settings and applications that we do not consider. This can include large utility-scale storage that can charge and discharge multiple consecutive hours on a MW scale.

Our results demonstrate that putting storage to multiple applications can introduce tradeoffs between the uses. Some storage applications can interfere with others giving subadditive values. In other cases, however, one application may increase the value of another use. Table 4 summarizes the value of the battery over the week that we simulate, and also annualizes this value on a \$/kWh-year basis. This annualized value computes the total value of the battery over a year, assuming that the week we examine is representative of the year, and normalizes the value based on the 11.2 kWh storage capacity of the battery. Table 4 also shows that distributed energy storage can be valuable, provided the value can be captured. Some of the value streams, for instance arbitrage and regulation, are easy to capture and quantify since these services are priced in wholesale markets. Avoided outage and curtailment costs can be more difficult to quantify, however, since they are averted costs that the customer does not bear. The value of these applications is also sensitive to the value of lost load, which may vary between customers. Avoided outages and curtailments can also provide value to the utility, which we do not consider. Using a battery to avoid a customer outage can help a utility meet its reliability requirements, while distribution relief can reduce loading on distribution-level transformers, potentially extending their lives. These value streams could further increase the value of distributed storage. While the design of contracts and incentive mechanisms that capture such value and allocate it to the owner is an important issue, it is beyond the scope of our work.

Table 4 Weekly and annualized value of battery under different cases.

	One-Week Value [$]	Annualized Value [$/kWh-year]
Arbitrage-Only	2.03	9.45
Backup Energy	2.05	9.54
Regulation	25.52	118.83
Distribution Relief	44.52	207.26

A Formulation of MIP

The formulation of the MIP, with the piecewise-linear approximation of the SDP value function, is given by:

$$\max_{a,x,n,q,y} C_t(S_t,a_t)+\gamma \sum_{j\in J_{t+1}} \text{Prob}\left\{\tilde{W}^j_{t+1}|W_t\right\}\cdot \sum_{m=1}^{M_{t+1}-1} \tilde{\sigma}^m_{t+1}(\tilde{W}^j_{t+1})\cdot q^{m,j}_{t+1};$$

$$\text{s.t.}\quad x^j_{t+1}=x_t+\eta^c(e^c_t+\tilde{\delta}^{d,j}_{t+1}k^d_t-n^{d,j}_{t+1})-(e^d_t+e^l_t+\tilde{\delta}^{u,j}_{t+1}k^u_t-n^{u,j}_{t+1})/\eta^d, \quad (15)$$
$$\forall\, j\in J_{t+1};$$
$$\underline{R}\le x^j_{t+1}\le \overline{R},\quad \forall\, j\in J_{t+1}; \quad (16)$$
$$0\le e^c_t+k^d_t\le \overline{P}^b; \quad (17)$$
$$0\le e^d_t+e^l_t+k^u_t\le \overline{P}^b; \quad (18)$$
$$-\overline{P}^h(1-I_t)\le l_t-e^l_t+e^c_t-e^d_t-k^u_t; \quad (19)$$
$$l_t-e^l_t+e^c_t-e^d_t+k^d_t\le \overline{P}^h(1-I_t); \quad (20)$$
$$n^{u,j}_{t+1}=\max\left\{0,\tilde{\delta}^{u,j}_{t+1}k^u_t-\eta^d(x_t-\underline{R})+e^d_t+e^l_t-e^c_t\right\},\quad \forall\, j\in J_{t+1}; \quad (21)$$
$$n^{d,j}_{t+1}=\max\left\{0,\tilde{\delta}^{d,j}_{t+1}k^d_t-(\overline{R}-x_t)/\eta^c-e^d_t-e^l_t+e^c_t\right\},\quad \forall\, j\in J_{t+1}; \quad (22)$$
$$l_t\le D_t; \quad (23)$$
$$e^d_t,e^c_t,k^d_t,k^u_t=0,\ \text{if } I_t=1; \quad (24)$$
$$e^d_t,e^c_t,e^l_t,l_t,k^u_t,k^d_t\ge 0; \quad (25)$$
$$n^{u,j}_{t+1},n^{d,j}_{t+1}\ge 0,\quad \forall\, j\in J_{t+1}; \quad (26)$$
$$\sum_{m=1}^{M_{t+1}-1} q^{m,j}_{t+1}=x^j_{t+1},\quad \forall\, j\in J_{t+1}; \quad (27)$$
$$0\le q^{m,j}_{t+1}\le y^{m,j}_{t+1}\cdot(\tilde{x}^{m+1}_t-\tilde{x}^m_t),\quad \forall\, j\in J_{t+1}, m=1,\dots,M_{t+1}-1; \quad (28)$$
$$y^{m,j}_{t+1}\le y^{m+1,j}_{t+1}\in\{0,1\},\quad \forall\, j\in J_{t+1}, m=1,\dots,M_{t+1}-1. \quad (29)$$

Constraints (15) through (26) are the same set of constraints in the original SDP, except that (15), (16), (21), (22), and (26) must hold for each possible realization of W^j_{t+1}. This is also why the hour-$(t+1)$ storage level is indexed by j, even though we do not restrict it to any discretization. Rather, x^j_{t+1} represents the hour-$(t+1)$ storage level resulting from the decisions chosen if the exogenous state, $\tilde{W}^j_{t+1}$, is realized, as defined by (15).

We linearize the max operators in constraints (21) and (22) by representing the right-hand-side of the equalities as piecewise-linear functions. Specifically, define $\underline{\Upsilon}^j_{t+1}$ and $\overline{\Upsilon}^j_{t+1}$ as the minimum and maximum value that the term $\tilde{\delta}^{u,j}_{t+1}k^u_t-\eta^d(x_t-\underline{R})+e^d_t+e^l_t-e^c_t$ can

take. We can then define two sets of auxiliary variables, $z_{t+1}^{\zeta,j}$ and $v_{t+1}^{\zeta,j}$, where $\zeta \in \{1,2\}$ and replace constraints (21) with:

$$n_{t+1}^{u,j} = \overline{\Upsilon}_{t+1}^{j} z_{t+1}^{2,j},$$

and add the constraints:

$$\begin{aligned}
&\sum_{\zeta=1}^{2} z_{t+1}^{\zeta,j} = \tilde{\delta}_{t+1}^{u,j} k_t^u - \eta^d (x_t - \underline{R}) + e_t^d + e_t^l - e_t^c, \quad \forall\, j \in J_{t+1};\\
&0 \le z_{t+1}^{1,j} \le -v_{t+1}^{1,j} \cdot \underline{\Upsilon}_{t+1}^{j}, \quad \forall\, j \in J_{t+1};\\
&0 \le z_{t+1}^{2,j} \le v_{t+1}^{2,j} \cdot \overline{\Upsilon}_{t+1}^{j}, \quad \forall\, j \in J_{t+1};\\
&v_{t+1}^{1,j} \le v_{t+1}^{2,j} \in \{0,1\}, \quad \forall\, j \in J_{t+1}.
\end{aligned}$$

Constraints (22) are linearized in the same manner. Constraints (27) through (29) define the piecewise-linear approximation of the cost-to-go function.

B Formulation of Deterministic Linear Program

The formulation of the deterministic linear program, which is used for the backcasting heuristic and the sample path average upper bound, is given by:

$$\begin{aligned}
\max_{a,x}\ & \sum_{\tau=t}^{T} \gamma^{\tau-t} C_\tau(S_\tau, a_\tau);\\
\text{s.t.}\ & x_{\tau+1} = x_\tau + \eta^c (e_\tau^c + \delta_{\tau+1}^d k_\tau^d - n_{\tau+1}^d) - (e_\tau^d + e_\tau^l + \delta_{\tau+1}^u k_\tau^u - n_{\tau+1}^u)/\eta^d,\\
&\qquad \forall\, \tau = t,\dots,T;\\
& \underline{R} \le x_\tau \le \overline{R}, \quad \forall\, \tau = t,\dots,T;\\
& 0 \le e_\tau^c + k_\tau^d \le \overline{P}^b, \quad \forall\, \tau = t,\dots,T;\\
& 0 \le e_\tau^d + e_\tau^l + k_\tau^u \le \overline{P}^b, \quad \forall\, \tau = t,\dots,T;\\
& -\overline{P}^h (1 - I_\tau) \le l_\tau - e_\tau^l + e_\tau^c - e_\tau^d - k_\tau^u, \quad \forall\, \tau = t,\dots,T;\\
& l_\tau - e_\tau^l + e_\tau^c - e_\tau^d + k_\tau^d \le \overline{P}^h (1 - I_\tau), \quad \forall\, \tau = t,\dots,T; \qquad (30)\\
& n_{\tau+1}^u = \max\left\{0, \delta_{\tau+1}^u k_\tau^u - \eta^d (x_\tau - \underline{R}) + e_\tau^d + e_\tau^l - e_\tau^c\right\}, \quad \forall\, \tau = t,\dots,T; \qquad (31)\\
& n_{\tau+1}^d = \max\left\{0, \delta_{\tau+1}^d k_\tau^d - (\overline{R} - x_\tau)/\eta^c - e_\tau^d - e_\tau^l + e_\tau^c\right\}, \quad \forall\, \tau = t,\dots,T; \qquad (32)\\
& l_\tau \le D_\tau, \quad \forall\, \tau = t,\dots,T;\\
& e_\tau^d, e_\tau^c, k_\tau^d, k_\tau^u = 0, \text{ if } I_\tau = 1, \quad \forall\, \tau = t,\dots,T;\\
& e_\tau^d, e_\tau^c, e_\tau^l, l_\tau, k_\tau^u, k_\tau^d, n_{\tau+1}^u, n_{\tau+1}^d \ge 0, \quad \forall\, \tau = t,\dots,T.
\end{aligned}$$

The constraints of the LP are the same as those of the SDP, except that they are only enforced for a single sample path of exogenous random variable realizations, since the LP is deterministic. The max operators in constraints (31) and (32) are linearized in the same manner used for the MIP outlined in Appendix A.

C Formulation of Two-Stage Stochastic Program

We present the deterministic equivalent of the two-stage stochastic program used to generate the SAA upper bounds. We define J as the number of leaves in the two-stage scenario tree

underlying the problem. A state variable with the superscript j and the subscript t denotes the value of that variable in hour t if the sample path leading to the j^{th} leaf of the scenario tree is realized. Similarly, a decision variable with the superscript j and the subscript t denotes the hour-t action taken if the sample path leading to the j^{th} leaf of the scenario tree is realized. The formulation of the stochastic program is given by:

$$\max_{a,x} \ \frac{1}{J}\sum_{j=1}^{J}\sum_{t=1}^{T}\gamma^{t-1}C_t(S_t^j, a_t^j);$$

$$\text{s.t. } x_{t+1}^j = x_t^j + \eta^c(e_t^{c,j} + \delta_{t+1}^{d,j}k_t^{d,j} - n_{t+1}^{d,j}) - (e_t^{d,j} + e_t^{l,j} + \delta_{t+1}^{u,j}k_t^{u,j} - n_{t+1}^{u,j})/\eta^d,$$
$$\forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$\underline{R} \le x_t^j \le \overline{R}, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$0 \le e_t^{c,j} + k_t^{d,j} \le \overline{P}^b, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$0 \le e_t^{d,j} + e_t^{l,j} + k_t^{u,j} \le \overline{P}^b, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$-\overline{P}^h(1 - I_t^j) \le l_t^j - e_t^{l,j} + e_t^{c,j} - e_t^{d,j} - k_t^{u,j}, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$l_t^j - e_t^{l,j} + e_t^{c,j} - e_t^{d,j} + k_t^{d,j} \le \overline{P}^h(1 - I_t^j), \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$n_{t+1}^{u,j} = \max\left\{0, \delta_{t+1}^{u,j}k_t^{u,j} - \eta^d(x_t^j - \underline{R}) + e_t^{d,j} + e_t^{l,j} - e_t^{c,j}\right\}, \tag{33}$$
$$\forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$n_{t+1}^{d,j} = \max\left\{0, \delta_{t+1}^{d,j}k_t^{d,j} - (\overline{R} - x_t^j)/\eta^c - e_t^{d,j} - e_t^{l,j} + e_t^{c,j}\right\}, \tag{34}$$
$$\forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$l_t^j \le D_t^j, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$e_t^{d,j}, e_t^{c,j}, k_t^{d,j}, k_t^{u,j} = 0, \text{ if } I_t^j = 1, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$e_t^{d,j}, e_t^{c,j}, e_t^{l,j}, l_t^j, k_t^{u,j}, k_t^{d,j}, n_{t+1}^{u,j}, n_{t+1}^{d,j} \ge 0, \quad \forall\ j = 1,\dots,J,\ t = 1,\dots,T;$$

$$e_1^{d,j} = e_1^{d,j'}, \quad \forall\ j, j' = 1,\dots,J; \tag{35}$$

$$e_1^{c,j} = e_1^{c,j'}, \quad \forall\ j, j' = 1,\dots,J; \tag{36}$$

$$e_1^{l,j} = e_1^{l,j'}, \quad \forall\ j, j' = 1,\dots,J; \tag{37}$$

$$l_1^j = l_1^{j'}, \quad \forall\ j, j' = 1,\dots,J; \tag{38}$$

$$k_1^{u,j} = k_1^{u,j'}, \quad \forall\ j, j' = 1,\dots,J; \tag{39}$$

$$k_1^{d,j} = k_1^{d,j'}, \quad \forall\ j, j' = 1,\dots,J; \tag{40}$$

$$x_1^j = \underline{R}, \quad \forall\ j = 1,\dots,J. \tag{41}$$

The stochastic program includes the same set of constraints in the original SDP. The max operators in constraints (33) and (34) are represented as piecewise-linear functions in the same manner used for the MIP outlined in Appendix A. Constraints (35) through (40) are nonanticipativity constraints, which force the stage-1 decisions in hour 1 to be the same for all J leaves of the scenario tree. Constraints (41) initialize the battery's starting storage level to be empty.

Acknowledgements The authors thank Warren Powell, SMART@CAR members, Suvrajeet Sen, Armin Sorooshian, the editors, and reviewers for helpful comments and suggestions. This work was supported by the SMART@CAR consortium, and by an allocation of computing time from the Ohio Supercomputer Center. Any opinions and conclusions expressed in this paper are those of the authors and do not necessarily represent those of SMART@CAR members.

References

1. Black, M., Strbac, G.: Value of Bulk Energy Storage for Managing Wind Power Fluctuations. IEEE Transactions on Energy Conversion **22**, 197–205 (2007)
2. Brown, R.E., Gupta, S., Christie, R.D., Venkata, S.S., Fletcher, R.H.: Distribution system reliability assessment using hierarchical Markov modeling. IEEE Transactions on Power Delivery **11**, 1929–1934 (1996)
3. Butler, P.C., Iannucci, J., Eyer, J.M.: Innovative Business Cases For Energy Storage In a Restructured Electricity Marketplace. Tech. Rep. SAND2003-0362, Sandia National Laboratories (2003)
4. van Casteren, J.F.L., Bollen, M.H.J., Schmieg, M.E.: Reliability assessment in electrical power systems: the Weibull-Markov stochastic model. IEEE Transactions on Industry Applications **36**, 911–915 (2000)
5. Deane, J.P., Ó Gallachóir, B.P., McKeogh, E.J.: Techno-economic review of existing and new pumped hydro energy storage plant. Renewable and Sustainable Energy Reviews **14**, 1293–1302 (2010)
6. Denholm, P., Kulcinski, G.L., Holloway, T.: Emissions and energy efficiency assessment of baseload wind energy systems. Environmental Science and Technology **39**, 1903–1911 (2005)
7. Denholm, P., Sioshansi, R.: The value of compressed air energy storage with wind in transmission-constrained electric power systems. Energy Policy **37**, 3149–3158 (2009)
8. DeSieno, C.F., Stine, L.L.: A Probability Method for Determining the Reliability of Electric Power Systems. IEEE Transactions on Reliability **R-14**, 30–35 (1965)
9. Drury, E., Denholm, P., Sioshansi, R.: The Value of Compressed Air Energy Storage in Energy and Reserve Markets. Energy **36**, 4959–4973 (2011)
10. Assessment of energy storage systems suitable for use by electric utilities. Tech. Rep. EPRI-EM-264, Electric Power Research Institute, Palo Alto, CA, USA (1976)
11. EPRI-DOE Handbook for Energy Storage for Transmission and Distribution Applications. Tech. Rep. 1001834, Electric Power Research Institute and the U.S. Department of Energy, Palo Alto, CA and Washington, DC, USA (2003)
12. Espinosa, J.R.: Implementation and integration of air conditioner, cycling at southern california edison. IEEE Transactions on Power Systems **2**, 792–798 (1987)
13. Eyer, J.M., Corey, G.: Energy Storage for the Electricity Grid: Benefits and Market Potential Assessment Guide. Tech. Rep. SAND2010-0815, Sandia National Laboratories (2010)
14. Eyer, J.M., Iannucci, J.J., Corey, G.P.: Energy Storage Benefits and Market Analysis Handbook. Tech. Rep. SAND2004-6177, Sandia National Laboratories (2004)
15. García-González, J., de la Muela, R.M.R., Santos, L.M., González, A.M.: Stochastic Joint Optimization of Wind Generation and Pumped-Storage Units in an Electricity Market. IEEE Transactions on Power Systems **23**, 460–468 (2008)
16. Graves, F., Jenkin, T., Murphy, D.: Opportunities for Electricity Storage in Deregulating Markets. The Electricity Journal **12**, 46–56 (1999)
17. Greenblatt, J.B., Succar, S., Denkenberger, D.C., Williams, R.H., Socolow, R.H.: Baseload wind energy: modeling the competition between gas turbines and compressed air energy storage for supplemental generation. Energy Policy **35**, 1474–1492 (2007)
18. Joo, S.K., Kim, J.C., Liu, C.C.: Empirical Analysis of the Impact of 2003 Blackout on Security Values of U.S. Utilities and Electrical Equipment Manufacturing Firms. IEEE Transactions on Power Systems **22**, 1012–1018 (2007)
19. Kariuki, K.K., Allan, R.N.: Evaluation of reliability worth and value of lost load. IEE Proceedings—Generation, Transmission, and Distribution **143**, 171–180 (1996)
20. Kempton, W., Tomić, J.: Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. Journal of Power Sources **144**, 268–279 (2005)
21. Knittel, C.R., Roberts, M.R.: An empirical examination of restructured electricity prices. Energy Economics **27**, 791–817 (2005)
22. Midlam-Mohler, S., Ewing, S., Marano, V., Guezennec, Y., Rizzoni, G.: PHEV Fleet Data Collection and Analysis. In: 2009 IEEE Vehicle Power and Propulsion Conference, pp. 1205–1210. Institute of Electrical and Electronics Engineers, Dearborn, MI, United States (2009)

23. Mohseni, P., Stevie, R.G.: Electric vehicles: Holy grail or Fool's gold. In: Power & Energy Society General Meeting, 2009, pp. 1–5. Institute of Electrical and Electronics Engineers, Calgary, AB (2009)
24. Mokrian, P., Stephen, M.: A Stochastic Programming Framework for the Valuation of Electricity Storage. In: 26th USAEE/IAEE North American Conference. International Association for Energy Economics, Ann Arbor, MI (2006)
25. Nourai, A.: Installation of the First Distributed Energy Storage System (DESS) at American Electric Power (AEP). Tech. Rep. SAND2007-3580, Sandia National Laboratories (2007)
26. Paatero, J.V., Lund, P.D.: Effect of energy storage on variations in wind power. Wind Energy **8**, 421–441 (2005)
27. Powell, W.B.: Approximate Dynamic Programming: Solving the Curses of Dimensionality. Wiley-Interscience, Hoboken, New Jersey (2007)
28. Seppälä, A.: Statistical distribution of customer load profiles. In: Proceedings of 1995 International Conference on Energy Management and Power Delivery. Institute of Electrical and Electronics Engineers (1995)
29. Shapiro, A.: Inference of statistical bounds for multistage stochastic programming problems. Mathematical Methods of Operations Research **58**, 57–68 (2003)
30. Shrestha, G.B., Songbo, Q.: Statistical Characterization of Electricity Price in Competitive Power Markets. In: 2010 IEEE 11th International Conference on Probabilistic Methods Applied to Power Systems (PMAPS). Institute of Electrical and Electronics Engineers, Singapore (2010)
31. Sioshansi, R.: Emissions Impacts of Wind and Energy Storage in a Market Environment. Environmental Science and Technology **45**, 10,728–10,735 (2011)
32. Sioshansi, R.: Increasing the Value of Wind with Energy Storage. The Energy Journal **32**, 1–30 (2011)
33. Sioshansi, R., Denholm, P.: Benefits of Colocating Concentrating Solar Power and Wind. IEEE Transactions on Sustainable Energy **4**, 877–885 (2013)
34. Sioshansi, R., Denholm, P., Jenkin, T.: A Comparative Analysis of the Value of Pure and Hybrid Electricity Storage. Energy Economics **33**, 56–66 (2011)
35. Sioshansi, R., Denholm, P., Jenkin, T., Weiss, J.: Estimating the Value of Electricity Storage in PJM: Arbitrage and Some Welfare Effects. Energy Economics **31**, 269–277 (2009)
36. Succar, S., Williams, R.H.: Compressed Air Energy Storage Theory, Resources, And Applications For Wind Power. Tech. rep., Princeton Environmental Institute (2008)
37. Tomić, J., Kempton, W.: Using fleets of electric-drive vehicles for grid support. Journal of Power Sources **168**, 459–468 (2007)
38. Tulpule, P., Marano, V., Rizzoni, G.: Energy management for plug-in hybrid electric vehicles using equivalent consumption minimisation strategy. International Journal of Electric and Hybrid Vehicles **2**, 329–350 (2010)
39. Walawalkar, R., Apt, J., Mancini, R.: Economics of electric energy storage for energy arbitrage and regulation in New York. Energy Policy **35**, 2558–2568 (2007)