

Optimal residential battery storage operations using robust data-driven dynamic programming

Nan Zhang, *Student Member, IEEE*, Benjamin D. Leibowicz, Grani A. Hanusanto

Abstract—In this paper, we consider the problem of operating a battery storage unit in a home with a rooftop solar photovoltaic (PV) system so as to minimize expected long-run electricity costs under uncertain electricity usage, PV generation, and electricity prices. Solving this dynamic program using standard techniques is computationally burdensome, and is often complicated by the difficulty of estimating conditional distributions from sparse data. To overcome these challenges, we implement a data-driven dynamic programming (DDP) algorithm that uses historical data observations to generate empirical conditional distributions and approximate the cost-to-go function. Then, we formulate two robust data-driven dynamic programming (RDDP) algorithms that consider the worst-case expected cost over a set of conditional distributions centered at the empirical distribution, and within a given Chi-square or Wasserstein distance, respectively. We test our algorithms using data from homes with rooftop PV in Austin, Texas. Numerical results reveal that DDP and RDDP outperform common existing methods with acceptable computational effort. Finally, we show that implementation of these superior operational algorithms significantly raises the break-even battery cost under which a homeowner is incentivized to invest in a residential battery rather than participate in a feed-in tariff or net energy metering program.

Index Terms—Battery, energy storage, solar PV, robust optimization, dynamic programming, stochastic control

I. INTRODUCTION

DISTRIBUTED renewable energy technologies are playing an increasingly important role in electricity generation due to benefits such as energy cost savings, reduced carbon emissions, and greater customer autonomy over energy choices [1]. Among distributed renewable generation options, solar photovoltaic (PV) systems have been widely deployed at the residential level [2], and residential PV capacity in the U.S. is projected to grow at a 10–15% annual rate between 2018 and 2022 [3]. However, unlike traditional, dispatchable generation facilities, the power output of a PV system is uncertain. Therefore, energy storage technologies that can be operated cost-effectively will be crucial for mitigating intermittency issues and enabling the continued expansion of PV generation.

This paper focuses on developing and testing algorithms for efficiently operating residential distributed energy systems that couple rooftop PV with battery storage (PV-battery systems). While residential PV ownership has grown rapidly, the capital cost of battery storage remains high. Clearly, formulating a more effective operational strategy for PV-battery systems will

The authors are with the Graduate Program in Operations Research and Industrial Engineering, The University of Texas at Austin, Austin, TX 78712, USA (email: nzhang17@utexas.edu; bleibowicz@utexas.edu; grani.hanusanto@utexas.edu).

significantly improve their generation usage efficiency and reduce electricity bills, which in turn will accelerate market uptake of PV-battery systems and amplify their cost and emissions reduction benefits.

We consider an individual household with a self-owned PV-battery system already installed. In each period, the household consumes a random quantity of electricity. Part of the demand is supplied from the PV system directly or by discharging the battery. The rest is obtained from the grid with a stochastic, time-varying electricity price. Our goal is to minimize the electricity bill by optimizing the stochastic control of battery operations under uncertain electricity usage, PV generation, and electricity prices.

The literature on energy storage operations is quite large. In [4], the authors developed an hourly-discretized optimization algorithm to find the optimal daily operational control of a wind-hydro power plant. Two-stage stochastic programming models were formulated in [5] and [6] to maximize profit by jointly optimizing wind generation and pumped hydroelectric storage. In [7], the authors used particle swarm optimization to maximize customers' net benefits by scheduling their distributed energy resources. In [8], researchers proposed an infinite-horizon stochastic control model to minimize the imbalance between available renewable power and load with small storage capacity. They represented the imbalance as a Laplace distributed process. The authors of [9] optimized demand-side management in a smart grid as a noncooperative game, and proved the existence of Nash equilibria.

A number of previous studies have applied dynamic programming to energy storage problems. In [10], an adaptive dynamic programming algorithm was constructed for grid management to deal with uncertain market conditions and customer behavior. In [11], a dynamic programming model with a priori forecasting was used to improve the performance of a hybrid system featuring wind and solar PV generation coupled with compressed air energy storage. A stochastic dynamic programming model was developed in [12] to co-optimize a distributed battery storage unit that provides multiple services which compete for its capacity. In [13], researchers studied an electricity storage problem with inelastic demand, and constructed an algorithm using approximate dynamic programming with a Markov transition matrix trained by real-world data. Online algorithms based on Lyapunov optimization with thresholds were proposed in [14] and [15] to optimize storage device scheduling, and these methods achieved asymptotic optimality. In [16], the author proposed Lyapunov optimization to optimize load scheduling and energy storage control simultaneously. In [17], a finite-horizon dynamic programming

model with an online stochastic algorithm combined with regularization was designed for storage utilization in a smart grid. In [18], the authors established a threshold structure for the optimal energy storage operational policy, which performed efficiently in a setting with random electricity prices.

In this paper, we formulate the problem in a data-driven manner and develop two robust data-driven dynamic programming algorithms to optimize PV-battery system operations. Our paper makes the following novel contributions.

- 1) To the best of our knowledge, we are the first researchers to apply data-driven dynamic programming (DDP) to the problem of optimizing the operations of residential energy systems with rooftop PV and battery storage. Our model directly leverages observed historical data, in contrast to most other dynamic programming approaches that use given or fitted distributions. This property gives our method an advantage for dealing with the residential battery storage operations problem, which has multiple stages, high-dimensional uncertainty, and often limited historical data that provide little or no prior knowledge about the true parameter distributions.
- 2) We construct two robust data-driven dynamic programming (RDDP) algorithms designed to help the DDP approach achieve better performance in out-of-sample circumstances. One algorithm uses Chi-square distance, and the other uses Wasserstein distance, to define the confidence set of distributions over which robust optimization minimizes the worst-case cost. Notably, our use of Wasserstein distance within a DDP framework is the first such attempt in the literature. Numerical results show that our RDDP algorithms perform significantly better than a state-of-the-art heuristic and approximate dynamic programming, with acceptable computational effort. Our RDDP scheme incorporating Wasserstein distance performs nearly as well as that with Chi-square distance, while substantially reducing computation time.
- 3) We analyze whether a household with rooftop PV should invest in a residential battery storage unit, or participate in a feed-in tariff or net metering program. By maximizing the value of a battery, implementation of our RDDP algorithms raises the break-even battery cost under which a household would find it optimal to install a battery. Therefore, our results quantify the impact that superior operational efficiency achieved through the RDDP control algorithms can have on accelerating the adoption of residential battery systems as their costs decline. We find that the impact is large.

The rest of the paper proceeds as follows. Section II introduces the generic dynamic programming model for battery storage operations, and then Section III outlines our DDP formulation. Section IV describes how we incorporate robustness to form the RDDP model. In Section V, we conduct numerical experiments using real-world data to compare the performance of our algorithms against previously proposed methods, and to investigate some factors that affect performance. In Section VI, we analyze the implications of our algorithms for the break-even battery cost under which a household would be

better off investing in a residential battery than participating in a feed-in tariff or net metering program. We conclude in Section VII with a summary of our most important findings and contributions.

II. PROBLEM FORMULATION

In this section, we describe the operation of a residential PV-battery system over T time periods and formulate a mathematical optimization model for the dynamic decision-making problem.

The residential unit is equipped with a PV panel that may supply additional energy to satisfy the household consumption and reduce the amount purchased from the grid. There is also a battery that can be utilized to store any excess energy for use at later times, with imperfect efficiency. Note that power flow between the home and grid is not explicitly modeled.

The operation of the residential PV-battery system takes place as follows. At the beginning of period t , there are s_t units of energy stored in the battery. Throughout the period, the PV panel generates w_t units of energy while the household consumes u_t units. These quantities may vary stochastically over time. Any unmet demand will be satisfied by discharging x_t units from the battery and by purchasing g_t units from the grid at a time-varying, per-unit price π_t . The system then transitions to the next time period $t + 1$ with a new battery storage level s_{t+1} , and the process is repeated until the terminal time T .

We now formalize the mathematical optimization model that represents the dynamic decision-making problem. To this end, we define by $\xi_t := (\pi_t, u_t, w_t)$ the triplet of *exogenous* parameters consisting of the grid electricity price, the household energy consumption, and the electricity generated by the PV panel in period t , respectively. Let $\xi_t = (\xi_1, \dots, \xi_t)$ denote the vector of historical exogenous parameters up to and including time t . The values of these parameters are independent of the decisions made by the system operator (i.e., the household). In contrast, the battery level s_t depends on the decisions made prior to time t and is often referred to as the *endogenous* state. The system operator's objective is to determine a sequence of purchasing and battery operation policies that minimize the total expected cost $\mathbb{E}[\sum_t \pi_t g_t]$ over the entire planning horizon T , where the decisions g_t and x_t are adapted to the state (s_t, ξ_t) , for all $t = 1, \dots, T$. This formulation gives rise to a stochastic optimal control problem whose parameters and decision variables are delineated in Table I.

The stochastic optimal control problem can be solved via a dynamic programming procedure, as follows. For any fixed battery storage level s_t and exogenous state ξ_t , the cost-to-go at time t is given by the optimal value of the optimization problem

TABLE I
PARAMETERS AND DECISION VARIABLES

Deterministic parameters:

ρ_c	charging efficiency
ρ_d	discharging efficiency
ρ_s	storage efficiency
S_{\max}	battery energy storage capacity
S_0	storage level at the beginning of the cycle
C	battery charging capacity per period
D	battery discharging capacity per period

Stochastic parameters:

π_t	grid electricity price in period t
u_t	electricity usage in period t
w_t	electricity generated by the PV panel in period t
ξ_t	triplet of exogenous parameters in period t , $\xi_t := (\pi_t, u_t, w_t)$

Decision variables:

g_t	electricity purchased from the grid in period t
x_t^d	electricity discharged from the battery in period t
x_t^c	electricity charged to the battery in period t
s_t	electricity stored in the battery at the beginning of period t

$$V_t(s_t, \xi_t) = \min \pi_t g_t + \mathbb{E}[V_{t+1}(s_{t+1}, \xi_{t+1}) | \xi_t] \quad (1a)$$

$$\text{s. t. } w_t + x_t^d - x_t^c + g_t \geq u_t \quad (1b)$$

$$g_t \in \mathbb{R}_+ \quad (1c)$$

$$x_t^c \in [0, C], x_t^d \in [0, D] \quad (1d)$$

$$s_{t+1} \in [0, S_{\max}] \quad (1e)$$

$$s_{t+1} \leq \rho_s s_t - \frac{x_t^d}{\rho_d} + \rho_c x_t^c \quad (1f)$$

Constraint (1b) stipulates that the demand u_t must be satisfied via electricity supplied by the PV panel w_t , the battery $x_t \equiv x_t^d - x_t^c$, and the grid g_t . (1c) restricts the power drawn from the grid to be non-negative (i.e., selling electricity back to the grid is not allowed). (1d) enforces maximum charge and discharge rates for the battery, and (1e) specifies its maximum energy storage capacity. (1f) models the evolution of the storage level over one time period, including losses due to imperfect storage, charging, and discharging efficiencies.

The objective function (1a) is defined as the sum of the purchase cost for time period t and the conditional expectation of the cost-to-go at time $t + 1$. By Bellman's principle of optimality, this objective function coincides with the expected cost over the time periods t, \dots, T . Thus, the total expected cost over the entire planning horizon is given by $V_1(s_1, \xi_1)$, which is obtained by solving the problem (1) backwards for $t = T, \dots, 1$, with a terminal condition $V_{T+1} \equiv 0$.

The dynamic programming procedure outlined above solves the stochastic optimal control problem exactly. The scheme is convex in the sense that i) for every fixed s_t and ξ_t , the optimization problem (1) is convex, and ii) for every fixed ξ_t , the cost-to-go $V_t(s_t, \xi_t)$ is a convex function in s_t . Despite these appealing properties, the scheme suffers

from two major shortcomings. First, it is computationally challenging due to the requirement of solving the problem (1) for the continuum of all states (s_t, ξ_t) . Also, evaluating the conditional expectation exactly is generally intractable as it involves multidimensional integration. Second, distributional knowledge about the exogenous parameters is typically incomplete. The decision maker solely has at his or her disposal a sequence of historical trajectories $\{\xi_T\}_{i=1}^N$ that can be utilized to infer the conditional distribution of ξ_{t+1} , and to estimate the conditional expectation in (1a). In the following section, we develop a data-driven scheme that addresses these two shortcomings.

III. DATA-DRIVEN DYNAMIC PROGRAMMING

In the DDP formulation, the conditional expectation in (1a) is replaced with its empirical estimate given by

$$\sum_{i=1}^N \hat{p}_{ti}(\xi_t) V_{t+1}(s_{t+1}, \xi_{t+1}^i). \quad (2)$$

The sample conditional probabilities $(\hat{p}_{ti}(\xi_t))_{i=1}^N$ in (2) are determined via the nearest-neighbor learning algorithm [19], which assigns positive mass only to the K closest observations to the reference point ξ_t , and simply neglects others. The nearest-neighbor learning algorithm enables us to select “good” trajectories automatically and improves the efficiency of the algorithm. The algorithm also utilizes a Gaussian smoother $\mathcal{S}(y) = \exp(\|y\|^2/2)/\sqrt{2\pi}$ so that data points nearer to ξ_t have larger weights. Specifically, we have

$$\hat{p}_{ti}(\xi_t) = \begin{cases} \frac{\mathcal{S}(\xi_t - \xi_t^i)}{\sum_{j \in \mathbb{N}(\xi_t, K_t)} \mathcal{S}(\xi_t - \xi_t^j)} & \text{if } i \in \mathbb{N}(\xi_t, K_t) \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where $\mathbb{N}(\xi_t, K_t)$ denotes the set of indices of the K_t closest data points to ξ_t . Since the distribution of candidate points in $\{\xi_t\}_{i=1}^N$ varies significantly with t , instead of a fixed value, we set a dynamic K_t depending on the dispersion of data for each t . In other words, K_t is determined by a given similarity threshold θ as

$$K_t = \min \left\{ k : \frac{\sum_{j \in \mathbb{N}(\xi_t, k)} \mathcal{S}(\xi_t - \xi_t^j)}{\sum_{j \in \mathbb{N}(\xi_t, N)} \mathcal{S}(\xi_t - \xi_t^j)} \geq \theta \right\}. \quad (4)$$

Replacing the conditional expectation in (1a) with the estimate (2) yields the formulation

$$\begin{aligned} \hat{V}_t(s_t, \xi_t) &= \min \pi_t g_t + \sum_{i=1}^N \hat{p}_{ti}(\xi_t) \hat{V}_{t+1}(s_{t+1}, \xi_{t+1}^i) \\ \text{s. t. } & g_t \in \mathbb{R}_+ \\ & x_t^c \in [0, C], x_t^d \in [0, D] \\ & s_{t+1} \in [0, S_{\max}] \\ & w_t + x_t^d - x_t^c + g_t \geq u_t \\ & s_{t+1} \leq \rho_s s_t - \frac{x_t^d}{\rho_d} + \rho_c x_t^c. \end{aligned}$$

The DDP scheme solves this problem backwards for $t = T, \dots, 1$ to arrive at an approximation $\hat{V}_1(s_1, \xi_1)$ of the true total expected cost $V_1(s_1, \xi_1)$, i.e., $\hat{V}_1(s_1, \xi_1) \approx V_1(s_1, \xi_1)$. This

scheme was originally studied in [20], where the conditional expectation is approximated via the Nadaraya-Watson kernel regression [21], [22]. In this paper, we provide an extension by incorporating nearest-neighbor learning, which enables a more efficient implementation when the data points are heavily clustered.

The DDP scheme is attractive because it is asymptotically consistent, meaning that the true total expected cost is recovered as the data size grows. It also results in a significant improvement in computational tractability as it only requires evaluating the cost-to-go function at the historical data points $\{\xi_t^i\}_{i=1}^N$. To further alleviate the intractability of evaluating the cost-to-go function for all endogenous states $s_t \in [0, S_{\max}]$, we discretize the state space into m points $\{0, \frac{1}{m-1}S_{\max}, \frac{2}{m-1}S_{\max}, \dots, S_{\max}\}$ and evaluate the function only at these points. The cost-to-go function is then approximated by a convex piecewise linear function where the interpolation points are given by the discretized points. Since the endogenous state is univariate and the cost-to-go function is convex in s_t , this piecewise linear architecture does not severely impact the computational complexity or the approximation quality of the DDP scheme.

IV. ROBUST DATA-DRIVEN DYNAMIC PROGRAMMING

When the historical data are sparse, the DDP scheme is known to generate an optimistically biased solution that overfits the given data points and performs poorly in out-of-sample tests. To mitigate this overfitting effect, we reinforce the scheme with tools from robust optimization. To this end, we relax the assumption that the empirical conditional expectation (2) constitutes a precise estimate of the true conditional expectation in (1a), and construct an *uncertainty set* of plausible conditional probabilities that are *close* to the nominal one in (3):

$$\mathcal{P}_\epsilon^\Phi(\hat{\mathbf{p}}_t) = \left\{ \mathbf{p} \in \mathbb{R}_+^N : \sum_{i=1}^N p_i = 1, \Phi(\mathbf{p}, \hat{\mathbf{p}}_t) \leq \epsilon \right\}. \quad (5)$$

Here, Φ is a prescribed metric that measures the distance between two probability distributions, while ϵ is a parameter that controls the degree of robustness.

The proposed RDDP scheme optimizes in view of the most adverse outcome from within the uncertainty set (5). Thus, in the formulation, the conditional expectation is replaced with the worst-case estimate

$$\max_{\mathbf{p} \in \mathcal{P}_\epsilon^\Phi(\hat{\mathbf{p}}_t)} \sum_{i=1}^N p_i V_{t+1}(s_{t+1}, \xi_{t+1}^i).$$

Solutions to the resulting min-max problem are more conservative and thereby less vulnerable to estimation errors in the empirical conditional probabilities. As with its DDP counterpart, the RDDP scheme is asymptotically consistent under an appropriate scaling of the robustness parameter ϵ . Furthermore, by the pointwise supremum property of convex functions, one can establish that the scheme is convex and also amenable to the same piecewise linear approximation architecture described in Section III. We now discuss two alternative methods to construct the distance function Φ that lead to tractable reformulations.

A. The Chi-square Distance

A natural way to define distances between two probability distributions is via the ϕ -divergences [23], [24]. In this paper, we adopt a class of ϕ -divergences specified by the Chi-square (χ^2) distance, where the distance between the probability vector \mathbf{p} and a reference vector $\hat{\mathbf{p}}$ is given by $\mathcal{C}(\mathbf{p}, \hat{\mathbf{p}}) = \sum_{i=1}^N (p_i - \hat{p}_i)^2 / p_i$. Here, the uncertainty set $\mathcal{P}_\epsilon^\mathcal{C}(\hat{\mathbf{p}}_t)$ corresponds to the set of probability vectors \mathbf{p} whose distance to $\hat{\mathbf{p}}_t$ is no greater than the prescribed level ϵ with respect to the χ^2 distance. Using this uncertainty set, the emerging robust problem is amenable to a tractable conic programming reformulation.

Proposition 1 (Theorem 4.1 in [20]): The RDDP formulation with the uncertainty set $\mathcal{P}_\epsilon^\mathcal{C}(\hat{\mathbf{p}}_t)$ is equivalent to the second-order cone program

$$\begin{aligned} \hat{V}_t^\mathcal{C}(s_t, \xi_t) = \min & \pi_t g_t + 2\lambda + \epsilon\lambda - \mu + \sum_{i=1}^N \hat{p}_{ti}(\xi_t) y_i \\ \text{s. t.} & g_t \in \mathbb{R}_+ \\ & x_t^c \in [0, C], x_t^d \in [0, D] \\ & s_{t+1} \in [0, S_{\max}] \\ & \mu \in \mathbb{R}, \lambda \in \mathbb{R}_+, \mathbf{z}, \mathbf{y} \in \mathbb{R}^N \\ & w_t + x_t^d - x_t^c + g_t \geq u_t \\ & s_{t+1} \leq \rho_s s_t - \frac{x_t^d}{\rho_d} + \rho_c x_t^c \\ & \hat{V}_t^\mathcal{C}(s_{t+1}, \xi_{t+1}^i) \leq z_i \quad \forall i \\ & z_i + \mu \leq \lambda \quad \forall i \\ & \|(y_i, z_i + \mu)\| \leq 2\lambda - z_i - \mu \quad \forall i. \end{aligned}$$

B. The Wasserstein Distance

The Wasserstein distance has recently become popular for determining robust solutions to one- and two-stage decision problems under uncertainty [25]. In this paper, we incorporate the Wasserstein distance into our RDDP scheme to compute robust solutions to the multi-stage battery operation problem. Formally, the Wasserstein distance between two distributions \mathbf{p} and $\hat{\mathbf{p}}$ supported on the same discrete set $\{\xi^1, \dots, \xi^N\}$ is defined as the optimal value of the linear program

$$\begin{aligned} \mathcal{W}(\mathbf{p}, \hat{\mathbf{p}}) = \min & \sum_{i,j \in [N]} \Pi_{ij} \|\xi^i - \xi^j\| \\ \text{s. t.} & \mathbf{\Pi} \in [0, 1]^{N \times N} \\ & \sum_{i,j \in [N]} \Pi_{ij} = 1 \\ & \sum_{j \in N} \Pi_{ij} = p_i \quad \forall i, \sum_{i \in N} \Pi_{ij} = \hat{p}_j \quad \forall j. \end{aligned} \quad (6)$$

The decision variable $\mathbf{\Pi} := (\Pi_{ij})_{i,j \in [N]}$ in (6) defines a joint probability distribution on the product space $\Xi \times \Xi$. The constraint system on the last line of (6) indicates that the joint distribution has marginals \mathbf{p} and $\hat{\mathbf{p}}$, respectively. The objective of the problem is to minimize the total cost of transporting a probability mass described by \mathbf{p} to one described by $\hat{\mathbf{p}}$. Here, the norm $\|\xi^i - \xi^j\|$ specifies the cost of transporting a Dirac point mass from ξ^i to ξ^j , while the decision variable Π_{ij} encodes the transportation plan.

Using the uncertainty set $\mathcal{P}_\epsilon^W(\hat{p}_t)$, the emerging robust problem is amenable to a more tractable linear programming reformulation.

Proposition 2: The RDDP formulation with the uncertainty set $\mathcal{P}_\epsilon^W(\hat{p}_t)$ is equivalent to the linear program

$$\begin{aligned}
 V_t^W(s_t, \xi_t) = \min \quad & \pi_t g_t + \lambda \epsilon - \mu + \sum_{i=1}^N \hat{p}_{ti}(\xi_t) y_i \\
 \text{s. t.} \quad & g_t \in \mathbb{R}_+ \\
 & x_t^c \in [0, C], x_t^d \in [0, D] \\
 & s_{t+1} \in [0, S_{\max}] \\
 & \mu \in \mathbb{R}, \lambda \in \mathbb{R}_+, \mathbf{z}, \mathbf{y} \in \mathbb{R}^N \\
 & w_t + x_t^d - x_t^c + g_t \geq u_t \\
 & s_{t+1} \leq \rho_s s_t - \frac{x_t^d}{\rho_d} + \rho_c x_t^c \\
 & \hat{V}_t^W(s_{t+1}, \xi_{t+1}^i) + \mu \leq z_i \quad \forall i \\
 & z_i - \lambda \|\xi_t^i - \xi_t^j\| \leq y_j \quad \forall i, j.
 \end{aligned}$$

Table II summarizes the numbers of variables and constraints in the Chi-square RDDP (CRDDP), Wasserstein RDDP (WRDDP), and DDP formulations.

TABLE II
NUMBERS OF CONSTRAINTS AND VARIABLES IN THE FORMULATIONS

Formulation	CRDDP	WRDDP	DDP
Variables	$\mathcal{O}(N)$	$\mathcal{O}(N)$	$\mathcal{O}(N)$
Linear Constraints	$\mathcal{O}(N)$	$\mathcal{O}(N^2)$	$\mathcal{O}(N)$
Second Order Cone Constraints	$\mathcal{O}(N)$	0	0

V. ALGORITHM PERFORMANCE COMPARISON

In this section, we conduct numerical experiments using real-world data to compare the performance of our DDP and RDDP algorithms against previously proposed schemes for battery storage operations, and to investigate the factors that most meaningfully affect performance.

A. Data Sources

We use empirical data from ten homes in the Mueller neighborhood of Austin, Texas that have been collected and provided by Pecan Street, Inc. [26]. These homes all have rooftop solar PV panels and detailed household-level electricity usage monitoring equipment, but do not have battery storage systems. The data include the hourly electricity consumption and PV generation of each household in 2016 and 2017. We also obtain the time-varying, wholesale electricity prices in Austin during those years from the Electric Reliability Council of Texas (ERCOT) [27]. Prices are reported for 15-minute intervals, and we average these prices over each hour. We implicitly assume that for each household, the data from 2016 and the data from 2017 follow the same distribution, and thus we can use the former to train our models and the latter to test them. All data are separated by season so that, for example, application of our algorithms to summer battery operations only uses historical training data from the previous summer. We set the cycle length to $T = 24$ hours and the number of trajectories to $N = 90$ for each season. For battery parameters, we set the storage capacity to $S_{\max} = 10$ kWh, the maximum charge and discharge rates to $C = D = S_{\max}/2$, and the

efficiency coefficients to $\rho_c = \rho_d = 0.99$ and $\rho_s = 1$ [13], [15]. We set the initial storage level to $S_0 = S_{\max}/2$ and assume there are no operating costs for the battery.

B. Numerical Results

We conduct experiments on the ten selected homes to compare the performances of our algorithms as well as various benchmarks. Previous work [14] showed that the optimal operational policy can be characterized by two thresholds. That is, if the net demand is higher than the upper threshold, discharge the battery to satisfy such demand; if the net demand is smaller than the lower threshold, charge the battery greedily; if the demand lies between them, do nothing. We compare our algorithms to one previously proposed heuristic based on such a policy.

A well-defined heuristic for the storage problem is the Threshold-Based Approximation (TBA) algorithm, constructed in [13]. TBA, as a kind of model predictive control method, uses the expectation of uncertainty data to obtain the optimal storage level s_t in each period t , then sets both thresholds and uses fixed operational policies as described above. We select TBA as a benchmark for comparison because TBA follows the simple guidelines of a threshold-based optimal control policy, with demonstrated good performance and fast solution speed relative to other similar heuristics [13].

We also implement an approximate dynamic programming (ADP) method as a benchmark, which is more complex than TBA. In the ADP, we discretize the uncertainties $\{\xi_t = \pi_t, u_t - w_t\}_{t=1}^T$ with δ possible discrete values each and model their evolution as a two-dimensional Markov chain. We assume that the storage level is discrete, within the set $\{0, \frac{1}{m-1}S_{\max}, \frac{2}{m-1}S_{\max}, \dots, S_{\max}\}$. Since the evolution of the storage level is uniquely determined by charging and discharging actions, the actions in one period can be viewed as the difference between the current storage level and the level in the next period, accounting for efficiency losses. Thus, the cardinality of the action space in each period is exactly m . We train the $\delta^2 \times \delta^2$ transition probability matrices for each period using historical observations. We set $\delta = 10$ to achieve balance between avoiding sparsity in transition probability matrices and maintaining the validity of historical data.

As additional benchmarks, we include the case where no battery storage system is installed (No Battery), and the case where the household has perfect information (PI). In the former case, the household must always purchase its net demand from the grid, and excess PV generation is wasted, as often happens in reality. In the latter case, the household starts off knowing the future values of all parameters with certainty, and optimizes deterministically. The PI objective value serves as a lower bound on the costs achieved by our algorithms. In our RDDP algorithms, ϵ is determined via cross-validation. The key outcome metrics we focus on for comparison are the average daily electricity cost in 2017, and the 95th percentile of the daily cost distribution. These numerical results are presented in Table III.

In Table III, we see that our DDP algorithm achieves a lower average daily cost than TBA and ADP in all ten homes.

Metric	Algorithm	1	2	3	4	5	6	7	8	9	10
Average Daily Cost	CRDDP	2.641	3.550	7.638	2.865	1.242	4.664	2.376	2.646	4.406	7.938
	WRDDP	2.655	3.643	7.683	2.888	1.242	4.680	2.493	2.656	4.417	7.972
	DDP	2.766	3.690	7.790	2.960	1.352	4.956	2.494	2.860	4.530	8.082
	ADP	3.028	4.353	7.953	3.517	1.756	5.126	3.228	3.028	4.606	8.092
	TBA	3.435	5.269	8.323	4.106	1.842	5.727	3.576	3.435	4.845	8.556
	No Battery	4.054	5.590	9.599	4.918	2.067	6.606	4.009	4.054	5.540	9.599
95 th Percentile Daily Cost	PI	2.165	2.891	6.606	2.373	0.945	3.958	1.966	2.165	3.781	6.881
	CRDDP	6.052	4.780	14.062	6.433	3.241	9.429	5.689	6.052	7.809	14.507
	WRDDP	6.079	4.991	14.017	6.443	3.318	9.503	6.045	6.079	8.324	14.632
	DDP	6.206	5.140	14.370	6.617	3.551	10.014	6.054	6.207	8.451	14.721
	ADP	6.653	5.554	14.673	7.458	3.901	10.000	6.914	6.653	8.043	14.872
	TBA	7.154	7.334	14.660	8.970	4.749	9.996	7.401	7.154	8.389	15.070
No Battery	8.114	6.778	17.695	9.341	4.749	12.387	8.013	8.114	9.809	17.695	
	PI	4.726	3.998	11.563	5.444	2.392	8.376	4.404	4.726	6.322	11.893

TABLE III
NUMERICAL RESULTS COMPARING OUR ALGORITHMS AND VARIOUS BENCHMARKS IN TEN HOMES

Moreover, our two RDDP algorithms outperform DDP. The CRDDP scheme achieves slightly lower costs than its WRDDP counterpart, but as we will show below, the WRDDP variant requires significantly less computation time. The advantages of the RDDP schemes relative to DDP are generally more substantial for the 95th percentiles of the daily electricity costs than for their means. This highlights a major advantage of incorporating robustness, as it reduces the likelihood of realizing very high costs.

C. Factors Affecting Performance

In this subsection, we investigate the factors which appear to exert the strongest influence on the performance of our algorithms.

The main challenges for residential battery storage operations are the high degrees of uncertainty in electricity usage and PV generation. These uncertainties are very seasonally dependent. Therefore, we separated both training and test data by season. Even though the data may be automatically clustered in the training process if we set an appropriate similarity threshold θ in nearest-neighbor learning, the seasonal categorization can still help us by reducing the size of the training data.




Fig. 1. Algorithm performance comparison, by season.

Fig. 1 illustrates the algorithm performance comparison for one sample home in each of the four seasons. We see

that the performance advantages of our RDDP algorithms are greater in spring and summer than in the other two seasons. The reasoning behind this is two-fold. First, in Austin, PV output is higher relative to electricity consumption in the spring and summer. With more excess PV output, the operational strategy for the battery is more important. Second, the uncertain parameters exhibit less variability in the fall and winter. With lower variability, our approximation of the cost-to-go function is relatively accurate, leaving little room to improve performance by incorporating robustness.

In addition to external factors such as the season, there are several internal factors within our model which might have a substantial impact on the performance of the algorithms. By inspecting the model outputs, we find that the cost in the final hour of the battery operational cycle comprises a large fraction of the total cost. This is a result of our policy that forces the battery to return to its initial storage level at the end of the cycle. In practical operations, it would be more reasonable to set the battery operational cycle T to a much longer duration, such as one week, and re-solve the dynamic program at more frequent intervals so that each solution is never actually implemented through its time horizon. However, increasing T would necessitate more historical observations as sample data to train the empirical conditional distribution for each period, because the dependencies between data from different days are not strong enough. In addition, using a longer cycle duration would raise the computation time of the algorithm.

Another important factor is the battery storage capacity S_{\max} , whose impact on algorithm performance relative to the PI lower bound is plotted in Fig. 2. As expected, TBA and DDP perform worse relative to PI when the battery capacity is larger. With a larger battery, the cost of charging or discharging the battery at inappropriate times is exaggerated due to a wider action space. By contrast, the performances of our CRDDP and WRDDP algorithms relative to PI do not appear to decline with the battery capacity. This can be viewed as evidence that the robustness incorporated into our RDDP algorithms successfully mitigates the negative effects of poor forecasting.

Battery storage capacity has a significant effect on the utilization efficiency of PV output. We calculate the daily curtailed (i.e., wasted) energy using different algorithms and




Fig. 2. Impact of the battery storage capacity S_{\max} .




Fig. 3. Impact of the robustness parameter ϵ .

storage capacities. We observe that our RDDP schemes outperform TBA and ADP on reducing energy curtailment. Particularly, compared to the No Battery benchmark, the utilization efficiencies of TBA and ADP decline slightly with increasing storage capacity, while the RDDP algorithms tend to perform better. These results are consistent with the relative performance of the algorithms in terms of the cost minimization objective, since limiting wasted energy is aligned with the overall goal of reducing grid electricity purchases.

Hyperparameter settings also play an important role in our algorithms. The uncertainty set in RDDP, which is governed by robustness parameter ϵ , can be seen to represent the risk attitude of the household. A relatively smaller ϵ implies that the household is more confident about its knowledge of future uncertainty based on past observations. We conduct sensitivity analysis to explore how ϵ can affect the performance of the algorithms using data from a single home and season, which can help us find near-optimal ϵ values when our algorithms are applied in different contexts. As shown in Fig. 3, for CRDDP, ϵ is fairly easy to calibrate via cross-validation. For WRDDP, it may be hard to find optimal ϵ values since the robustness is also measured by the geometric properties of the sample data, which vary with time index t . A dynamic setting of $\epsilon(t)$ could be implemented with affordable time complexity.

Lastly, we investigate the impact of the similarity threshold θ . Table IV reports the electricity costs achieved using our algorithms as well as their corresponding computational time

consumption. We observe that time consumption increases sharply as θ approaches 1, with only slight improvements in performance. Although we cannot guarantee theoretical properties when $\theta < 1$, it would seem preferable to set $\theta \approx 0.99$ in order to balance performance against computation time.

TABLE IV
ALGORITHM COMPARISON FOR DIFFERENT θ VALUES

θ	Average Daily Cost (\$)			Time Consumption (s)		
	CRDDP	WRDDP	DDP	CRDDP	WRDDP	DDP
0.95	2.3604	2.3634	2.4513	386	92	69
0.97	2.3443	2.3419	2.4403	630	94	71
0.99	2.3385	2.3417	2.4402	902	110	73
0.995	2.3229	2.3429	2.4363	1172	123	77
1	2.3180	2.3393	2.4359	1504	911	119

In conclusion, we find that CRDDP offers the best performance, while WRDDP performs only slightly worse while allowing substantial reductions in computation time.

VI. CASE STUDY

In this section, we assess the extent to which the performance gains realized through implementation of our algorithms can affect the decision of an individual household to invest in a battery storage unit or participate in a feed-in tariff (FIT) or net energy metering (NEM) program. In recent years, many utilities have begun offering these programs to encourage greater adoption of residential PV systems.

A FIT program allows households to sell all the electricity generated by their own PV systems to the utility at a fixed price. Meanwhile, households must pay for their entire electricity usage – even that which is obtained from their own PV systems – at the standard utility electricity rates. FITs require two power meters to independently measure the inflow and outflow electricity of a household, allowing consumption and generation to be priced separately. A FIT program was first proposed in the U.S. in 1978, and subsequently adopted by more than 60 countries in the past 40 years [28].

Unlike with a FIT, in a NEM program, households are only charged for their net electricity consumption at standard utility electricity rates. When the PV system generates surplus electricity, the household can sell the excess back to the grid to get credits, which are deducted from the electricity bill. NEM only requires a single bidirectional meter to measure power flowing in both directions, which is easier to implement than the two meters needed for a FIT. NEM originated in the U.S. in 1980 and has since been implemented in at least 38 states [29]. NEM programs across the U.S. differ in terms of details such as the maximum monthly generation net-metered or the deduction policy for credits earned by the household.

FIT and NEM are the two most common sell-back programs for residential electricity customers, and there are large literatures analyzing their effects. In most settings, households are forbidden from pairing battery electricity storage with participation in these programs. Utilities clearly are not willing to let households earn profit simply by charging their batteries during periods of low electricity prices and discharging to the




Fig. 4. Annual electricity cost, inclusive of annualized battery capital cost for strategies featuring batteries, as a function of the per-kWh battery cost.

grid when prices are high. In some jurisdictions, like Massachusetts, pairing battery storage with a sell-back program is tentatively accepted but strictly constrained; for instance, the battery storage capacity is limited, or the battery is not allowed to be charged from the grid [30].

From the perspective of a household with rooftop PV panels, there are two ways to lower its electricity bill. Under typical, flat retail electricity rates, there is little room to reduce costs by using a battery, which means that participating in a sell-back program will likely dominate investing in a battery at high capital cost. However, under time-varying electricity prices, it might be economical to invest in a battery rather than participate in a sell-back program. The more effectively the battery is operated, the more likely investing in it is the optimal choice.

In the following analysis, we compare battery investment to virtual FIT and NEM programs in our Austin setting. For the FIT, we set the sell-back price at 26.8¢/kWh, based on the real sell-back price in Austin's Value of Solar Tariff program with a five-tiered price structure [31]. For the NEM program, we assume that credits earned by a household are obtained immediately upon sell-back at a discount on the real-time electricity price. We assume that these credits are worth 75% of the real-time prices.

We analyze a sample home with high PV output relative to electricity consumption. Our reasoning is three-fold. First, the relative performance gains realized by implementing our algorithms are higher for such homes. Second, this home with high PV output relative to electricity usage is actually fairly representative of our full sample of homes. Third, for homes with low PV generation, savings from battery storage are unlikely to cover the battery capital cost, so they have less incentive to install a storage system.

Fig. 4 shows how the annual electricity cost, inclusive of the annualized battery capital cost for those strategies that feature battery storage, varies with the per-kWh battery cost. The different lines represent the household's alternatives, ranging from No Battery, to participating in the FIT or NEM program, to investing in a battery and operating it using the TBA,

DDP, or CRDDP algorithm. The pink PI line indicates the lower bound for all policy-free results, but this bound is not applicable in the presence of a FIT or NEM policy. The dashed vertical line, included for reference, indicates the per-kWh capital cost of a Tesla Powerwall 2 battery unit [32], a popular residential storage system. This unit has a 13.5 kWh storage capacity and a capital cost of \$6,600, implying a cost of \$488/kWh. To incorporate the cost of this system into the annual electricity costs, we amortize it over a 15-year lifetime (same as the warranty duration) at an assumed 2% interest rate.

The vertical distances between the TBA, DDP, and CRDDP lines correspond to the annual cost savings realized through implementation of our algorithms. As shown, simply by using CRDDP as opposed to TBA, this household would save nearly \$200 per year in electricity costs with a battery. At the current battery cost represented by the dashed vertical line, the household is better off participating in the sell-back programs than investing in a battery. However, use of the CRDDP scheme shifts the break-even points between the battery alternatives and the sell-back programs to the right, meaning that the household will be incentivized to adopt a battery system sooner as the cost of battery storage continues to decline. Using the TBA, DDP, and CRDDP algorithms, the per-kWh battery cost would have to fall to \$56, \$111, and \$247, respectively, for battery investment to become superior to the FIT program. The analogous break-even battery costs for the NEM program are negative for TBA and DDP, and \$122 for CRDDP. Clearly, our CRDDP algorithm raises the break-even battery cost significantly, and its performance advantages could thus have a significant impact on accelerating residential battery uptake.

Lastly, we explore what happens if the household is allowed to combine a battery with the FIT program, assuming that the battery is operated according to our CRDDP algorithm. This appears as the gray line in Fig. 4. Note that the FIT is equivalent to the household selling its PV generation at a fixed price independent of its consumption. Therefore, we can analyze the combined scenario by simply running our model without any solar generation, then subtracting the revenue gained from selling the generation. Fig. 4 shows that the break-even battery cost of the combined scenario with respect to the FIT scenario is greater than the current battery cost of \$488/kWh. So, if battery investment were permitted under the FIT scheme, this household would find it optimal to install a battery at current cost rather than participate in the FIT program without one.

VII. CONCLUSION

In this paper, we developed a DDP model that minimizes long-run electricity costs by stochastically controlling the operation of a residential battery storage unit under uncertain electricity usage, PV generation, and grid electricity prices. Our approach leverages actual historical data observations, in contrast to most existing dynamic programming schemes that rely on given or fitted distributions. To mitigate overfitting and improve out-of-sample performance, we constructed two

RDDP algorithms which consider the worst-case expected cost over a confidence set of distributions centered at the empirical conditional distribution. The CRDDP variant employs the χ^2 distance to establish this set, while the WRDDP version uses the Wasserstein distance.

We tested our algorithms on real data from homes with rooftop PV in Austin, Texas, and compared their performance against a state-of-the-art heuristic and an approximate dynamic programming method. Numerical results show that our DDP approach outperforms these other schemes, and that the RDDP algorithms achieve even greater cost savings, with affordable computational complexity. CRDDP offers the best performance, but by using the Wasserstein distance, WRDDP performs only slightly worse while substantially reducing computation time.

We then explored how the choice of operational algorithm influences the decision problem of a household deciding whether to invest in a residential battery system or participate in a FIT or NEM sell-back program. By operating the battery using CRDDP instead of the heuristic, a household can save hundreds of dollars per year in electricity costs, which significantly raises the break-even battery cost under which battery investment is preferred to FIT or NEM participation. Therefore, the performance advantages enabled by our algorithms developed in this paper could meaningfully accelerate the market uptake of residential battery storage.

ACKNOWLEDGMENT

Grani A. Hanususanto is supported by the National Science Foundation grant no. 1752125.

REFERENCES

- [1] C. Schelly, "Residential solar electricity adoption: what motivates, and what matters? a case study of early adopters," *Energy Research & Social Science*, vol. 2, pp. 183–191, 2014.
- [2] P. Denholm, R. Margolis, B. Palmintier, C. Barrows, E. Ibanez, L. Bird, and J. Zuboy, *Methods for analyzing the benefits and costs of distributed photovoltaic generation to the US electric utility system*. National Renewable Energy Laboratory, 2014.
- [3] A. Perea, "Solar market insight report 2017 q4." <https://www.seia.org/research-resources/solar-market-insight-report-2017-q4>.
- [4] E. D. Castronovo and J. P. Lopes, "On the optimization of the daily operation of a wind-hydro power plant," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1599–1606, 2004.
- [5] P. Mokrian and M. Stephen, "A stochastic programming framework for the valuation of electricity storage," in *26th USAEE/IAEE North American Conference*, pp. 24–27, Citeseer, 2006.
- [6] J. Garcia-Gonzalez, R. M. R. de la Muela, L. M. Santos, and A. M. Gonzalez, "Stochastic joint optimization of wind generation and pumped-storage units in an electricity market," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 460–468, 2008.
- [7] M. A. A. Pedrasa, T. D. Spooner, and I. F. MacGill, "Coordinated scheduling of residential distributed energy resources to optimize smart home energy services," *IEEE Transactions on Smart Grid*, vol. 1, no. 2, pp. 134–143, 2010.
- [8] H.-I. Su and A. El Gamal, "Modeling and analysis of the role of energy storage for renewable integration: Power balancing," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4109–4117, 2013.
- [9] I. Atzeni, L. G. Ordóñez, G. Scutari, D. P. Palomar, and J. R. Fonollosa, "Demand-side management via distributed energy generation and storage optimization," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 866–876, 2013.
- [10] M. Boaro, D. Fuselli, F. De Angelis, D. Liu, Q. Wei, and F. Piazza, "Adaptive dynamic programming algorithm for renewable energy scheduling and battery management," *Cognitive Computation*, vol. 5, no. 2, pp. 264–277, 2013.
- [11] V. Marano, G. Rizzo, and F. A. Tiano, "Application of dynamic programming to the optimal management of a hybrid power plant with wind turbines, photovoltaic panels and compressed air energy storage," *Applied Energy*, vol. 97, pp. 849–859, 2012.
- [12] X. Xi, R. Sioshansi, and V. Marano, "A stochastic dynamic programming model for co-optimization of distributed energy storage," *Energy Systems*, vol. 5, no. 3, pp. 475–505, 2014.
- [13] S. Kwon, Y. Xu, and N. Gautam, "Meeting inelastic demand in systems with storage and renewable sources," *IEEE Transactions on Smart Grid*, vol. 8, no. 4, pp. 1619–1629, 2017.
- [14] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, "Optimal energy storage control policies for the smart power grid," in *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, pp. 475–480, IEEE, 2011.
- [15] S. Lakshminarayana, Y. Xu, H. V. Poor, and T. Q. Quek, "Cooperation of storage operation in a power network with renewable generation," *IEEE Transactions on Smart Grid*, vol. 7, no. 4, pp. 2108–2122, 2016.
- [16] T. Li and M. Dong, "Real-time residential-side joint energy storage management and load scheduling with renewable integration," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 283–298, 2016.
- [17] A. S. Bedi, M. W. Ahmad, S. Swapnil, K. Rajawat, S. Anand, et al., "Online algorithms for storage utilization under real-time pricing in smart grid," *International Journal of Electrical Power & Energy Systems*, vol. 101, pp. 50–59, 2018.
- [18] J. Jin and Y. Xu, "Optimal storage operation under demand charge," *IEEE Transactions on Power Systems*, vol. 32, no. 1, pp. 795–808, 2017.
- [19] L. Devroye, L. Györfi, and G. Lugosi, *A probabilistic theory of pattern recognition*, vol. 31. Springer Science & Business Media, 2013.
- [20] G. A. Hanususanto and D. Kuhn, "Robust data-driven dynamic programming," in *Advances in Neural Information Processing Systems*, pp. 827–835, 2013.
- [21] E. A. Nadaraya, "On estimating regression," *Theory of Probability & its Applications*, vol. 9, no. 1, pp. 141–142, 1964.
- [22] G. S. Watson, "Smooth regression analysis," *Sankhyā: The Indian Journal of Statistics, Series A*, vol. 26, no. 4, pp. 359–372, 1964.
- [23] F. Liese and I. Vajda, "On divergences and informations in statistics and information theory," *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4394–4412, 2006.
- [24] A. Ben-Tal, D. Den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen, "Robust solutions of optimization problems affected by uncertain probabilities," *Management Science*, vol. 59, no. 2, pp. 341–357, 2013.
- [25] R. Gao and A. J. Kleywegt, "Distributionally robust stochastic optimization with wasserstein distance," *arXiv preprint arXiv:1604.02199*, 2016.
- [26] "Pecan street inc. dataport." <http://www.pecanstreet.org>.
- [27] "The electric reliability council of texas." <http://www.ercot.com/mktinfo/prices>.
- [28] J. Huenteler, "International support for feed-in tariffs in developing countries: review and analysis of proposed mechanisms," *Renewable and Sustainable Energy Reviews*, vol. 39, pp. 857–873, 2014.
- [29] "Database of state incentives for renewables and efficiency, net metering (raleigh, nc: North carolina state university, 2017)." <http://programs.dsireusa.org/system/program>.
- [30] T. J. Calter and P. W. Mark, "Bill h.2852: An act relative to net metering, community shared solar and energy storage." <https://www.seia.org/research-resources/solar-market-insight-report-2017-q4>.
- [31] "Value of solar (vos) rate." <https://austinenergy.com/ae/rates/residential-rates/value-of-solar-rate>.
- [32] "Meet powerwall, your home battery." <https://www.tesla.com/powerwall>.



Nan Zhang received a B.S. in Industrial Engineering from Tsinghua University in Beijing, China in 2017. He is currently pursuing a Ph.D. in Operations Research and Industrial Engineering at The University of Texas at Austin. His research interests include simulation, scheduling, and energy system operations under uncertainty.



Benjamin D. Leibowicz earned a B.A. in Physics from Harvard University, and received both M.S. and Ph.D. degrees in Management Science and Engineering from Stanford University. Presently, he is an Assistant Professor in the Graduate Program in Operations Research and Industrial Engineering at The University of Texas at Austin. He is also affiliated with the Department of Mechanical Engineering, Graduate Program in Energy and Earth Resources, and Lyndon B. Johnson School of Public Affairs (by courtesy). His research interests include energy systems, energy policy, technological change, and sustainable cities.



Grani A. Hanusanto received an MSc degree in Financial Engineering from the National University of Singapore and a PhD degree in Operations Research from Imperial College London. He was a postdoctoral researcher at the College of Management of Technology at Ecole Polytechnique Federale de Lausanne. He is currently an Assistant Professor of Operations Research and Industrial Engineering at The University of Texas at Austin. His research focuses on the design and analysis of tractable solution schemes for decision-making problems under uncertainty, with applications in operations management, energy systems, machine learning, and data analytics.