

Data-Driven Distributionally Robust Control of Energy Storage to Manage Wind Power Fluctuations

Samantha Samuelson

Insoon Yang

Abstract—Energy storage is an important resource that can balance fluctuations in energy generation from renewable energy sources, such as wind, to increase their penetration. Many existing storage control methods require perfect information about the probability distribution of uncertainties. In practice, however, the distribution of renewable energy production is difficult to reliably estimate. To resolve this challenge, we develop a new storage operation method, based on the theory of distributionally robust stochastic control, which has the following advantages. First, our controller is robust against errors in the distribution of uncertainties such as power generated from a wind farm. Second, the proposed method is effective even with a small number of data samples. Third, the construction of our controller is computationally tractable due to the proposed duality-based dynamic programming method that converts infinite-dimensional minimax optimization problems into semi-infinite programs. The performance of the proposed method is demonstrated using data about energy production levels at wind farms in the Pennsylvania-Jersey-Maryland interconnection (PJM) area.

I. INTRODUCTION

To decarbonize the electric power grid, there have been growing efforts to utilize clean renewable energy sources. The utilization of wind and solar generation is challenging because these energy sources are uncertain, intermittent and non-dispatchable. One solution is energy storage, a method of storing excess energy which can later be used to compensate for unexpected supply shortages [1]. The rapidly decreasing cost of large-scale energy storage devices also makes them attractive resources that can provide flexibility to the power grids [2]. California, for example, plans to deploy 1325 MW of utility-scale energy storage by the year 2020 [3]. This accounts for 2% of the state's peak demand. In this paper, we consider the use of energy storage specifically to balance the fluctuations in energy availability from wind farms with limited information about the distribution of wind generation outputs.

Previous work has focused on using stochastic optimal control methods to optimally control energy storage devices to account for the variability of wind and solar energy generation. For example, the dynamic-programming approaches developed in [4], [5], [6], [7] assume that the true shape of the distribution of wind energy availability is known exactly. These assumptions are rarely met, and accurate distribution models are difficult to obtain, for a variety of reasons. Particularly when historical wind power data is

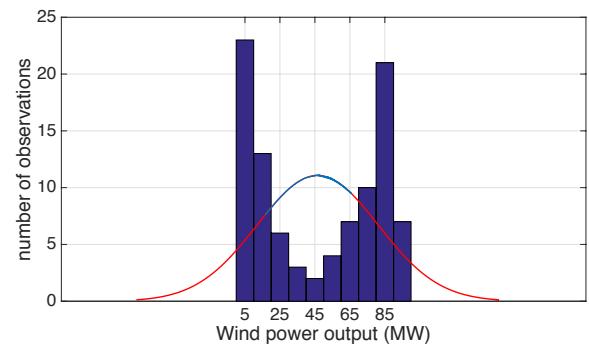


Fig. 1: Distribution of wind generation at 3 PM ($t = 15$) in January 2004-2006 with mean $\mu = 46.6$ and standard deviation $\sigma = 34.56$, generated from 48 observations. A scaled Gaussian with the same mean and variance is shown as well.

scarce, empirical estimations are likely to be highly inaccurate. Furthermore, a variable such as wind or solar energy is likely to vary over time, making even distribution estimates based on real data unreliable. Even when distributional data is available, attempting to infer the true distribution from data can often be challenging, and yield misleading results. For example, consider the data of wind energy shown in Fig. 1 for a particular hour at wind farms in the PJM area: the data distribution is clearly not Gaussian, and it would be extremely difficult to define the continuous probability density function which best describes this data. Nor are mean and variance alone enough to describe this behavior. While it is possible to define a discrete probability mass function which perfectly fits this set of data, to do so would certainly be overfitting and would contain inaccuracies.

To account for these difficulties, we propose a new stochastic control technique that is robust to uncertainty with regards to the true distribution of the random variable, in this case power produced by wind generators. This method is based on distributionally robust stochastic control, which minimizes the expected value of a given cost function without assuming the distribution takes a known shape, but instead that it is drawn from a known set of distributions (e.g., [8], [9], [10], [11], [12]). In essence, this approach formulates a control strategy which minimizes expected cost in the face of the worst-case distribution possible, given the information that is known about the distribution. It is worth mentioning that there exist storage control techniques which do not require the exact distribution of uncertainties [13], [14]. However, these approaches do not aim to design a

This work is supported in part by NSF CRII: CPS under CNS-1657100. S. Samuelson and I. Yang are with the Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA {sasamuel, insoonya}@usc.edu

controller that is robust against distributional errors in outputs of wind generators, unlike our method.

The proposed storage control method has the following advantages. First, the proposed controller is robust against distribution errors in data of uncertainties such as power produced by a wind farm. This is a direct consequence of the proposed minimax formulation of designing stochastic controllers. Second, our method can design an effective controller with a small number of data samples. To define admissible distributions, we employ a moment-based ambiguity set, which can be constructed from a small data set. We use historical wind data to compare the performance of the proposed distributionally robust control method with standard stochastic control methods using a limited number of wind data samples, and demonstrate the advantages of distributional robustness. Third, our duality-based dynamic programming method resolves computational challenges by reformulating infinite-dimensional minimax optimization problems as semi-infinite programs. This reformulation allows us to avoid discretization of the noise distribution. We also address the challenges of a non-convex cost function by optimizing two separate convex components.

The remainder of this paper is organized as follows. In Section II, we introduce the specifics of energy storage systems, including the state-of-charge dynamics and cost structure. A duality-based dynamic programming solution method is proposed in Section III. Section IV details the application of this control strategy to a real-world data set of wind energy, and the comparison to a standard stochastic controller.

II. PROBLEM SETUP

A. Energy Storage Model

Consider an energy storage system (ESS) such as a sodium sulfur (NaS) battery, lithium-ion battery, or compressed air energy storage. Let $x_t \in \mathcal{X} := [\underline{x}, \bar{x}]$ be the state of charge (SOC) in the ESS at stage t . The dynamics of the SOC can be written as

$$x_{t+1} = \eta(x_t + u_t), \quad t = 0, \dots, T-1, \quad (1)$$

where the control input $u_t \in \mathcal{U} := [\underline{u}, \bar{u}]$ represents the amount of useful energy to store and the coefficient η accounts for the losses that occur between two consecutive stages. We assume that $u_t \geq 0$ and $u_t \leq 0$ represent charging and discharging, respectively. Since $\underline{x} \leq x_t \leq \bar{x}$ and $\underline{u} \leq u_t \leq \bar{u}$, the charging/discharging control input must satisfy the constraint $u_t \in \mathbb{U}(x_t)$, where

$$\mathbb{U}(x_t) := \left[\max \left\{ \underline{u}, \frac{\bar{x}}{\eta} - x_t \right\}, \min \left\{ \bar{u}, \frac{\bar{x}}{\eta} - x_t \right\} \right]. \quad (2)$$

The amount of energy drawn from a bus to charge the ESS by $u_t \geq 0$ is greater than u_t because charging efficiency, denoted by α_c , is less than 1. On the other hand, discharging efficiency, denoted by α_d , is less than 1, and thus the amount of energy injected to the bus by discharging $-u_t \geq 0$ amount of energy from the ESS is less than $-u_t$. Specifically, the

total amount of energy drawn from the bus by the ESS is given by

$$h(u_t) := \frac{1}{\alpha_c} \max\{u_t, 0\} - \alpha_d \max\{-u_t, 0\}. \quad (3)$$

B. Ambiguity of Wind Power Distribution

Let $\{w_t\}_{t=0}^T$ be a stochastic process that represents wind power minus demand (in MW) at the bus connected with a controllable energy storage device of interest. Its distribution measure is denoted by μ_t . Unfortunately, in many cases of wind energy sources, the probability distribution of w_t cannot be reliably estimated due to the variability of wind generation outputs. In practice, for example, we may be able to estimate only the mean and variance of w_t with a good accuracy. To compensate for the variability of wind power using energy storage, we develop a new stochastic control method that is robust against errors in the estimated distribution of wind power. In this work, we use information about the mean and variance of w_t . This motivates us to introduce the following ambiguity set of probability distributions:

$$\mathbb{D}_t := \{\mu_t \in \mathcal{P}(\mathbb{R}) \mid \mu_t(W_t) = 1, \quad (4a)$$

$$|\mathbb{E}_{\mu_t}[w_t] - \mathbf{m}_t| \leq \mathbf{b}_t, \quad (4b)$$

$$\mathbb{E}_{\mu_t}[(w_t - \mathbf{m}_t)^2] \leq \mathbf{c}_t \sigma_t^2\}, \quad (4c)$$

where $\mathcal{P}(X)$ denotes the set of Borel probability measures on X . Constraint (4a) requires the distribution μ_t to be completely defined on support W_t . Constraints (4b) and (4c) constrain the mean and variance of the distribution μ_t to be within some confidence parameters \mathbf{b}_t and \mathbf{c}_t of the nominal mean and variance \mathbf{m}_t and σ_t^2 .

C. Distributionally Robust Energy Balancing

The total net demand at the bus taking into account the ESS is given by

$$\delta_t(u_t, w_t) := h(u_t) - w_t - s_t,$$

where s_t is the amount of power supplied by conventional (coal) generators to the bus. Its value is deterministic and often determined in a day-ahead electricity market. Considering an energy balancing problem, we set the cost function as

$$r_t(u_t, w_t) := p_t^+ \max\{\delta_t(u_t, w_t), 0\} + p_t^- \max\{-\delta_t(u_t, w_t), 0\},$$

where p_t^+ and p_t^- are the penalties for each unit of positive and negative energy imbalance at stage t , respectively, such that $p_t^+ \geq p_t^-$. Note that this stage-wise cost function is piecewise linear but non-convex, in general.

To minimize the worst-case penalties for energy imbalances with limited information about the distribution of wind generation output, we consider the following distributionally robust storage control problem:

$$\inf_{\pi \in \Pi} \sup_{\gamma \in \Gamma} \mathbb{E}^{\pi \cdot \gamma} \left[\sum_{t=0}^{T-1} r_t(u_t, w_t) + q(x_T) \right],$$

where $q : \mathbb{R} \rightarrow \mathbb{R}$ is a terminal cost function of interest. We assume that the control decision u_t is made after observing (x_t, w_t) . Thus, w_t is considered as a second state variable with stochastic evolution, where the future state w_{t+1} is drawn from some unknown distribution belonging to the ambiguity set \mathbb{D}_{t+1} . Here, the set Π of admissible control policies is given by $\Pi := \{\pi = (\pi_0, \dots, \pi_{T-1}) \mid u_t = \pi_t(x_t, w_t) \in \mathbb{U}(x_t)\}$. Similarly, the set Γ of admissible wind power distribution policies is given by $\Gamma := \{\gamma = (\gamma_0, \dots, \gamma_{T-1}) \mid \mu_{t+1} = \gamma_t(x_t, w_t) \in \mathbb{D}_{t+1}\}$. Note that the set \mathbb{D}_t encodes the ambiguity of wind power distribution. We also set $q(x_T) := C|x_T - x_{des}|$, where x_{des} is the desired final battery state, and C is a fixed penalty term, to penalize the deviation of the final SOC level from the desired range. This terminal penalty term is important because, for example, the storage device must continue to operate for the next day when $T = 24$ h. Note that the designed controller minimizes the worst-case total penalty no matter how the wind power distribution changes within the ambiguity set \mathbb{D}_t . Thus, the proposed control tool improves the robustness of the closed-loop system with respect to imperfect (empirical) distribution information about wind generation output.

III. DUALITY-BASED DYNAMIC PROGRAMMING

In order to determine the globally optimal storage control policy $\pi^* \in \Pi$, we use dynamic programming to solve this distributionally robust stochastic control problem.

A. Bellman Equation

We first define the value function, $v_t(\mathbf{x}, \mathbf{w})$ which simply denotes the minimum total cost of the system's continued evolution from time t to $T-1$, starting at state $(x_t, w_t) = (\mathbf{x}, \mathbf{w})$. The value function is defined as follows:

$$v_t(\mathbf{x}, \mathbf{w}) :=$$

$$\inf_{\pi \in \Pi} \sup_{\gamma \in \Gamma} \mathbb{E}^{\pi \cdot \gamma} \left[\sum_{s=t}^{T-1} r_s(u_s, w_s) + q(x_T) \mid (x_t, w_t) = (\mathbf{x}, \mathbf{w}) \right].$$

Thus, the value function gives the optimal, that is minimum, total cost in the face of the worst-case possible net demand (or wind) distribution in the ambiguity set \mathbb{D}_t . The optimal control strategy π is then the argmin of this term if it exists.¹

The value function can also be defined recursively as follows. For any given state (\mathbf{x}, \mathbf{w}) and control input \mathbf{u} at stage t , $v_{t+1}(\eta(\mathbf{x} + \mathbf{u}), w_{t+1})$ gives the minimum worst-case cost from stage $t+1$, using the storage dynamics given in (1). Thus, the cost from stage t , for any given state (\mathbf{x}, \mathbf{w}) and control input \mathbf{u} , can be expressed as $r_t(\mathbf{x}, \mathbf{u}) + v_{t+1}(\eta(\mathbf{x} + \mathbf{u}), w_{t+1})$. Essentially, the value at stage t is the sum of the stage-wise or running cost r_t at stage t , and the expected future optimal cost from stage $t+1$ that depends on the state resulting from the control action at stage t . This is the key idea of dynamic programming, expressed by the following

¹The conditions under which the distributionally robust control problem admits an optimal solution can be found in our previous work [11].

Bellman equation:

$$v_t(\mathbf{x}, \mathbf{w}) = \inf_{\mathbf{u} \in \mathbb{U}(\mathbf{x})} \sup_{\boldsymbol{\mu} \in \mathbb{D}_{t+1}} \mathbb{E}_{\boldsymbol{\mu}} [r_t(\mathbf{u}, \mathbf{w}) + v_{t+1}(\eta(\mathbf{x} + \mathbf{u}), w_{t+1})]. \quad (5)$$

This equation can be very difficult to solve numerically, due both to the scalability issues inherent to dynamic programming and to the fact that we are optimizing over the infinite-dimensional set \mathbb{D}_{t+1} of possible distributions. To solve this problem, we reformulate this problem as a single minimization problem over several variables by taking the dual of the inner maximization problem.

B. Dual Formulation

We can rewrite the inner maximization problem in the Bellman equation (5) more clearly by making the implicit constraints (4) of the optimization domain $\boldsymbol{\mu} \in \mathbb{D}_{t+1}$ explicit, as follows:

$$\begin{aligned} & \sup_{\boldsymbol{\mu} \in M_+(\mathbb{R})} \int_{W_{t+1}} r_t(\mathbf{u}, \mathbf{w}) + v_{t+1}(\eta(\mathbf{x} + \mathbf{u}), w_{t+1}) d\boldsymbol{\mu}(w_{t+1}) \\ \text{s.t. } & \mathbf{m}_{t+1} - \mathbf{b}_{t+1} \leq \int_{W_{t+1}} w_{t+1} d\boldsymbol{\mu}(w_{t+1}) \\ & \int_{W_{t+1}} w_{t+1} d\boldsymbol{\mu}(w_{t+1}) \leq \mathbf{m}_{t+1} + \mathbf{b}_{t+1} \\ & \int_{W_{t+1}} (w_{t+1} - \mathbf{m}_{t+1})^2 d\boldsymbol{\mu}(w_{t+1}) \leq \mathbf{c}_{t+1} \boldsymbol{\sigma}_{t+1}^2 \\ & \boldsymbol{\mu}(W_{t+1}) = 1. \end{aligned}$$

To find the dual of this problem, we first define a few constants and variables. Let $\mathbf{b}_{t+1} := \mathbf{b}_{t+1} - \mathbf{m}_{t+1}$ and $\mathbf{b}_{t+1} := \mathbf{b}_{t+1} + \mathbf{m}_{t+1}$. Let $\underline{\lambda} \in \mathbb{R}$ and $\bar{\lambda} \in \mathbb{R}$ be Lagrangian multipliers associated with the first two inequality constraints. Let $\Lambda \in \mathbb{R}$ be the Lagrangian multiplier associated with the second moment constraint. Finally, let $\nu \in \mathbb{R}$ be the Lagrangian multiplier associated with the final support constraint. We can then obtain the following dual optimization problem:

$$\begin{aligned} & \inf_{\underline{\lambda}, \bar{\lambda}, \Lambda, \nu} r_t(\mathbf{u}, \mathbf{w}) + \mathbf{b}_{t+1}^\top \underline{\lambda} + \bar{\mathbf{b}}_{t+1}^\top \bar{\lambda} + \mathbf{c}_{t+1} \boldsymbol{\sigma}_{t+1}^2 \Lambda + \nu \\ \text{s.t. } & (\bar{\lambda} - \underline{\lambda}) w_{t+1} + \Lambda (w_{t+1} - \mathbf{m}_{t+1})^2 + \nu \\ & \geq v_{t+1}(\eta(\mathbf{x} + \mathbf{u}), w_{t+1}) \quad \forall w_{t+1} \in W_{t+1} \\ & \underline{\lambda}, \bar{\lambda}, \Lambda \geq 0. \end{aligned}$$

This problem satisfies the assumptions in [11] under which there is no duality gap. Thus, we can substitute the dual minimization problem over finite-dimensional Lagrange multipliers with the original inner maximization problem in (5), to obtain the following dual Bellman equation:

$$\begin{aligned} & v_t(\mathbf{x}, \mathbf{w}) = \\ & \inf_{\mathbf{u}, \underline{\lambda}, \bar{\lambda}, \Lambda, \nu} r_t(\mathbf{u}, \mathbf{w}) + \mathbf{b}_{t+1}^\top \underline{\lambda} + \bar{\mathbf{b}}_{t+1}^\top \bar{\lambda} + \mathbf{c}_{t+1} \boldsymbol{\sigma}_{t+1}^2 \Lambda + \nu \\ \text{s.t. } & (\bar{\lambda} - \underline{\lambda}) w_{t+1} + \Lambda (w_{t+1} - \mathbf{m}_{t+1})^2 + \nu \\ & \geq v_{t+1}(\eta(\mathbf{x} + \mathbf{u}), w_{t+1}) \quad \forall w_{t+1} \in W_{t+1} \\ & \underline{\lambda}, \bar{\lambda}, \Lambda \geq 0, \quad \mathbf{u} \in \mathbb{U}(\mathbf{x}) \end{aligned} \quad (6)$$

with terminal condition $v_T(\mathbf{x}, \mathbf{w}) = q(\mathbf{x})$. Note that this is a semi-infinite program, which can be solved by existing convergent algorithms (see Section III-C). Due to the existence result in [11], for any $(t, \mathbf{x}, \mathbf{w})$ the optimal control action is the minimizer $\pi_t^*(\mathbf{x}, \mathbf{w})$ and we obtain the optimal Markov policy $\pi^* = \{\pi_t^*\}_{t=0}^{T-1}$.

C. Algorithm

We can numerically solve the dual Bellman equation by discretizing the state space. Initializing the value function at terminal stage T as $q(\mathbf{x})$ for each \mathbf{w} , we compute values $v_t(\mathbf{x}, \mathbf{w})$ recursively, by solving the optimization problem stated in (6) for each possible (discretized) value of \mathbf{x} and \mathbf{w} . Of course, while discretizing the state space allows us to find discrete values of $v_t(\mathbf{x}, \mathbf{w})$, we must obtain an expression for the function $v_t(\eta(\mathbf{x} + \mathbf{u}), w_{t+1})$ in order to solve for $v_{t-1}(\mathbf{x}, \mathbf{w})$. We resolve this challenge by using linear interpolation to create a set of 1-D continuous functions of \mathbf{x} passing through the values $v_t(\mathbf{x}^{(i)}, \mathbf{w}^{(j)})$, $i = 1, \dots, N_x$ for each $\mathbf{w}^{(j)}$, $j = 1, \dots, N_w$, where $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N_x)}$ and $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(N_w)}$ are pre-defined grid points.

An additional obstacle is presented by the cost function, which is not necessarily convex in \mathbf{u} . However, we circumvent this technicality by noting that the concavity is caused by a discontinuity in the first derivative at $\mathbf{u} = 0$.

Proposition 1. *We define two functions $r_t^+ : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ and $r_t^- : \mathbb{R}_- \times \mathbb{R} \rightarrow \mathbb{R}$ as*

$$\begin{aligned} r_t^+(\mathbf{u}, \mathbf{w}) &:= r_t(\mathbf{u}, \mathbf{w}) \quad \text{for } \mathbf{u} \geq 0 \text{ and } \mathbf{w} \in \mathbb{R} \\ r_t^-(\mathbf{u}, \mathbf{w}) &:= r_t(\mathbf{u}, \mathbf{w}) \quad \text{for } \mathbf{u} \leq 0 \text{ and } \mathbf{w} \in \mathbb{R}. \end{aligned}$$

Then, $r_t(\mathbf{u}, \mathbf{w}) = r_t^+(\mathbf{u}, \mathbf{w}) + r_t^-(\mathbf{u}, \mathbf{w})$ for all $(\mathbf{u}, \mathbf{w}) \in \mathbb{R} \times \mathbb{R}$ and the following convexity results hold:

- $r_t^+(\mathbf{u}, \mathbf{w})$ is convex in $\mathbf{u} \geq 0$ for each $\mathbf{w} \in \mathbb{R}$; and
- $r_t^-(\mathbf{u}, \mathbf{w})$ is convex in $\mathbf{u} \leq 0$ for each $\mathbf{w} \in \mathbb{R}$.

Proof. Recall that the total amount of energy drawn by the ESS is given in (3). For $\mathbf{u} \geq 0$, $h(\mathbf{u}) = \frac{1}{\alpha_c} \mathbf{u}$. We then note that $r_t^+(\mathbf{u}, \mathbf{w}) = p^+ \max\{\frac{1}{\alpha_c} \mathbf{u} - \mathbf{w} - s_t, 0\} + p^- \max\{-\frac{1}{\alpha_c} \mathbf{u} + \mathbf{w} + s_t, 0\}$ since $\mathbf{u} \geq 0$. Each element of the maximization is linear in \mathbf{u} . The point-wise maximum of convex functions is convex, and thus $r_t^+(\mathbf{u}, \mathbf{w})$ is convex in $\mathbf{u} \geq 0$.

We now turn our attention to $\mathbf{u} \leq 0$, where we have $h(\mathbf{u}) = \alpha_d \mathbf{u}$ due to (3). Thus, for $\mathbf{u} \leq 0$, we obtain that $r_t^-(\mathbf{u}, \mathbf{w}) = p^+ \max\{\alpha_d \mathbf{u} - \mathbf{w} - s_t, 0\} + p^- \max\{-\alpha_d \mathbf{u} + \mathbf{w} + s_t, 0\}$, which has the same structure as r_t^+ . Therefore, the previous argument holds for the convexity of $r_t^-(\mathbf{u}, \mathbf{w})$ in $\mathbf{u} \leq 0$. \square

An example of the cost function r_t and its subfunctions r_t^+ and r_t^- can be found in [15]. Inspired by the convexity result in Proposition 1, we split the optimization problem into two subproblems, assuming \mathbf{u} is positive or negative, and compare optimal cost to obtain the final solution. These steps are collected in Algorithm 1.

Note that for each time step and discretized state point, we have to solve the semi-infinite program (6). We use

Algorithm 1: Distributionally robust controller design

```

Initialize  $v_T(\mathbf{x}, \mathbf{w}) = q(\mathbf{x}) \quad \forall(\mathbf{x}, \mathbf{w})$ ;
discretize  $\mathbf{x} \rightarrow \mathbf{x}^{(i)}$ ,  $i = 1, \dots, N_x$ ;
discretize  $\mathbf{w} \rightarrow \mathbf{w}^{(j)}$ ,  $j = 1, \dots, N_w$ ;
for  $t = (T - 1)$  to 1 do
  create  $v_{t+1}(\mathbf{x}, \mathbf{w}^{(j)})$  using linear interpolation for
   $j = 1, \dots, N_w$ ;
  set time-dependent parameters  $\mathbf{b}_t, \mathbf{c}_t, \mathbf{m}_t, \sigma_t$ ;
  for  $\mathbf{x}^{(i)} = \mathbf{x}^{(1)}$  to  $\mathbf{x}^{(N_x)}$  do
    determine the constraint  $\mathbb{U}(\mathbf{x}^{(i)})$  of  $\mathbf{u}$  using (2);
    for  $\mathbf{w}^{(j)} = \mathbf{w}^{(1)}$  to  $\mathbf{w}^{(N_w)}$  do
      solve (6) for  $v_t^+$  with  $\operatorname{argmin} \mathbf{u}^* \geq 0$ ;
      solve (6) for  $v_t^-$  with  $\operatorname{argmin} \mathbf{u}^* \leq 0$ ;
      let  $v_t(\mathbf{x}^{(i)}, \mathbf{w}^{(j)}) = \min[v_t^+, v_t^-]$  and
      let  $\pi_t(\mathbf{x}^{(i)}, \mathbf{w}^{(j)})$  be the corresponding  $\mathbf{u}^*$ ;
    end for
  end for
end for

```

the convergent discretization method developed in [16] to solve the problem. However, one can employ other numerical methods such as primal and dual methods, homotopy methods, discretization methods, and exchange methods (see [17], [18] and the references therein).

IV. BALANCING WIND ENERGY

A. The Setup

We apply the proposed method to the problem of balancing wind energy by controlling a large-scale NaS battery. We start with data containing hourly energy production levels at 10 different wind farms in the PJM area over a period of three years. We focus on data from one wind farm during the months of January and March, each of which gives a set of 96 data points for each hour of the day. First we use equation (4) to construct an ambiguity set based on our observations so far of the random variable w_{t+1} , in this case wind level. When constructing a moment-based ambiguity set from historical data, we use empirical mean and variance for nominal mean and variance \mathbf{m}_t and σ_t , as well as minimum and maximum observed values of w_t to define the support W_t .² We also set $\mathbf{b}_t = 0$ and $\mathbf{c}_t = 1$. Once ambiguity sets are created, we use Algorithm 1 to find the optimal control policy $\pi_t^*(x_t, w_t)$ at each stage t .

We use the model described in Section II, with the following parameter values: $\alpha_c = \alpha_d = 0.85$, $p^+ = 100$, $p^- = -200$, $\eta = 0.97$, and $C = 500$. We also let $(\underline{x}, \bar{x}) = (5, 50)$ MWh, and $(\underline{u}, \bar{u}) = (-30, 30)$ MW.

²We assume that a portion of demand is balanced by conventional (coal) generators, and set unmet demand d_t to be the 25% quantile of the empirical distribution with $s_t = 0$. Therefore, we redefine the random variable w_t by shifting as follows: $w_t \leftarrow w_t - d_t$. Thus, at each time the unbalanced wind power w_t has some non-zero mean, standard deviation σ_t , and support $W_t = [w_{\min}(t), w_{\max}(t)] - d_t$.

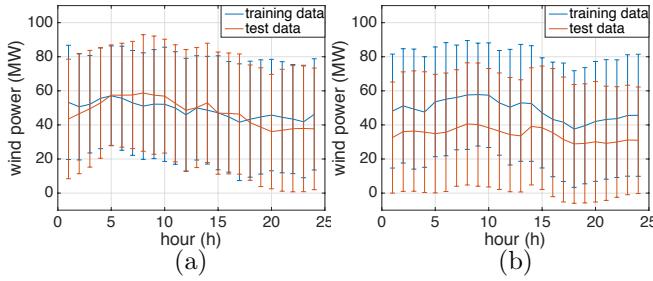


Fig. 2: The mean-variance plot of the training (blue) and test (red) data sets for the month of (a) January and (b) March.

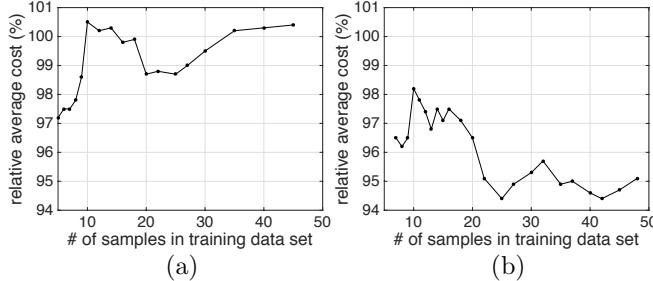


Fig. 3: The average cost incurred by the distributionally robust controller as a percentage of the average cost incurred by the standard controller: results with (a) the January data set and (b) the March data set. Performance averaged over 10,000 sample trajectories.

B. Comparison to a Standard Stochastic Controller

To evaluate the performance of our distributionally robust controller, we compare against a standard dynamic programming-based stochastic controller. We use empirical data to find the expected value of the cost function by formulating the observed data as a probability mass function, where each observed value is equally likely. We divide the available data at each time into two sets: a training set and a test set. The training data set is used to construct ambiguity sets and train the standard controller, while the test data set is used for evaluation. For these results we consistently use the last 48 data points as our test set, whose mean-variance plots are shown in Fig. 2.

We evaluate the performance of different controllers by simulating the evolution of the system based on the input of either controller. We initialize $x_0 = x_{des} = \frac{1}{2}\bar{x}$. At each time t , the variable w_{t+1} is drawn randomly from the corresponding test data set, by selecting one of the 48 datapoints, each with equal probability. We compare the average total cost incurred by each controller over 10,000 trials. We first examine the performance of the distributionally robust controller as compared to the size of the training data set for each month. As shown in Fig. 3 (a), in January for very small training data sets, the distributionally robust controller performs slightly better than the standard controller, earning a total average cost up to 3% smaller than the standard stochastic controller. However, as the size of the training data set grows, it provides a more accurate

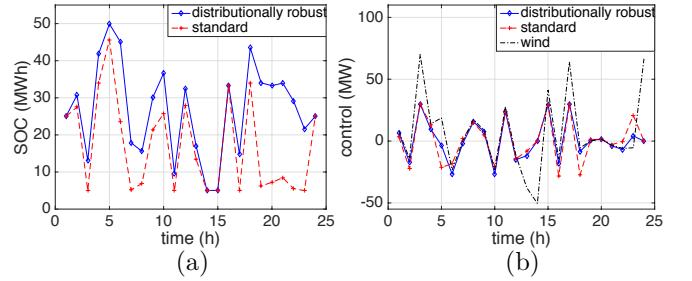


Fig. 4: (a) The state-of-charge trajectory under the distributionally robust controller (blue) and the standard stochastic controller (red). (b) The two control trajectories. The wind realization at each time is traced in black.

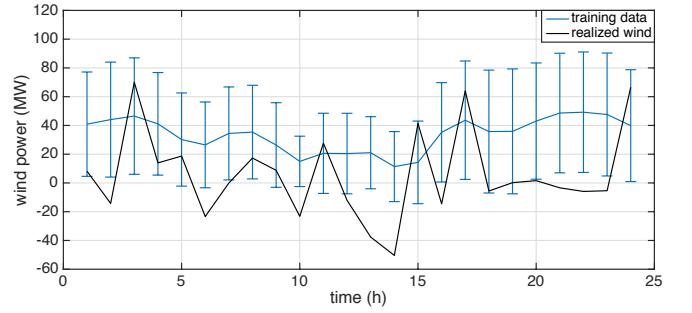


Fig. 5: The distribution of 8-sample training data set (blue) with a realization of wind power overlaid in black.

estimate of the true distribution of the test data set, and the distributionally robust controller loses its advantage. We observe a different behavior when evaluating the performance of the distributionally robust controller on data from March, as shown in Fig. 3 (b). This figure shows that the distributionally robust controller performs close to 6% better than the standard controller, and the relative performance does not deteriorate with the size of the training data set. To understand this difference, we examine the training and test data sets for each month, as shown in Fig. 2. From these mean-variance plots of the training and test data distributions, we can see that both data sets are fairly similar for January, as opposed to having remarkably different distributions for March. The distributionally robust controller demonstrates a superior performance in situations where the test data is not represented by the training data. This is the case for March, but not for January. It is important to note that while the distributionally robust controller offers little advantage for larger data sets in January, it does not exhibit a performance significantly inferior to the stochastic controller either. Faced with choosing a controller for a month with an unknown test data distribution, it is still safest to use the distributionally robust controller.

To understand why the distributionally robust controller performs better than the standard controller in most cases, we now focus on the case of a training data set consisting of 8 samples, from the month of March. The state and control trajectories for a particular realization of wind energy randomly sampled from the test data sets are shown in Fig. 4. On

this trial, the distributionally robust controller achieves a total cost that is 67.8% of the total cost incurred by the standard controller. Fig. 4 (b) shows the wind realization overlaid in black. From this, we can see how both controllers are primarily driven to match the realization of wind generation output w_t at each time step, while satisfying the terminal conditions, with roughly similar state trajectories. However, at times 2, 5, 16, and 18 h, we can see the distributionally robust controller accurately meet the wind realization while the standard controller discharges the battery more than necessary. In particular, at 18 h the battery SOC is close to saturation as shown in Fig. 4 (a). The training data set predicts the wind realization at the next time step will be large, with a mean of approximately 40 MW as shown in Fig. 5. The distributionally robust controller, which chooses a control value to match the wind at 18 h, retains a high SOC at 19 h, and thus the maximum allowed control value is approximately 15 MW. The standard controller discharges the battery by a greater amount at 18 h, resulting in a lower SOC. Thus, the standard controller has a higher maximum allowed control value at 19 h of 30 MW, allowing it to better match a hypothetical large wind power realization. Note that it is cheaper for the standard controller to overly discharge the battery at 18 h than to be unable to absorb excess wind at 19 h. Essentially, the standard controller is choosing to incur a small cost in the present to prevent a larger cost at the next time stage. However, this trade-off only makes sense if the controller is sure of the future cost. Hence, the distributionally robust controller simply matches wind power realization at 18 h, as the ambiguity set expresses less confidence in the prediction of high wind levels at 19 h. As shown in Fig. 5, the prediction of the training data set was incorrect, and wind energy at 19 h was low, resulting in a lower cost for the distributionally robust controller. We can observe from Fig. 2 that the training set consists of data with a higher mean than the test data set, and most realizations of wind energy will have a lower mean than the training data set. The distributionally robust controller is less willing to incur cost in the present to prevent possible future costs, being less certain of future wind estimates, as shown in the example of times 18 and 19 h. Thus, it has an advantage when the training data is sufficiently different than the test data.

V. CONCLUSIONS AND FUTURE WORK

We investigated a distributionally robust stochastic control method in the use of energy storage to balance fluctuations in wind power. Using the same sample data to design ambiguity sets and construct a standard stochastic controller, we find that the distributionally robust controller has a superior performance when training data is sparse (in the case of January), or when there is a fundamental difference in the distribution of training and test data sets (in the case of March). The distributionally robust controller has a more conservative response to predictions of extreme values of future wind energy. This clearly demonstrates the advantages of the distributionally robust controller for the application

of energy balancing, where our estimate of wind power distribution is likely to be inaccurate.

Future work involves expanding the distributionally robust controller to take advantage of the potential correlation of wind generation outputs across time. Possible future work also includes coordinating energy storage systems across the power grids, which results in an interesting network control problem. Another possible direction is to refine the construction of ambiguity sets. While it is intuitive to create a family of distributions by constraining their mean and variance to be within some bound, it may be useful to group similar distributions using Wasserstein distance or confidence intervals.

REFERENCES

- [1] A. Castillo and D. F. Gayme, "Grid-scale energy storage applications in renewable energy integration: A survey," *Energy Conversion and Management*, vol. 87, pp. 885–894, 2014.
- [2] B. Nykvist and M. Nilsson, "Rapidly falling costs of battery packs for electric vehicles," *Nature Climate Change*, vol. 5, pp. 329–332, 2015.
- [3] California Public Utilities Commission, "Decision adopting energy storage procurement framework and design program," 2013.
- [4] J. H. Kim and W. B. Powell, "Optimal energy commitments with storage and intermittent supply," *Operations Research*, vol. 59, no. 6, pp. 1347–1360, 2011.
- [5] 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.
- [6] P. Harsha and M. Dahleh, "Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy," *IEEE Transactions on Power Systems*, vol. 30, no. 3, pp. 1164–1181, 2015.
- [7] S. Lakshminarayana, Y. Xu, H. V. Poor, and T. Q. S. 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.
- [8] H. Xu and S. Mannor, "Distributionally robust Markov decision processes," *Mathematics of Operations Research*, vol. 37, no. 2, pp. 288–300, 2012.
- [9] P. Yu and H. Xu, "Distributionally robust counterpart in Markov decision processes," *IEEE Transactions on Automatic Control*, vol. 61, no. 9, pp. 2538–2543, 2016.
- [10] B. P. G. Van Parys, D. Kuhn, P. J. Goulart, and M. Mori, "Distributionally robust control of constrained stochastic systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2016.
- [11] I. Yang, "A dynamic game approach to distributionally robust safety specifications for stochastic systems," *arXiv:1701.06260 [math.OC]*, 2017.
- [12] —, "A convex optimization approach to distributionally robust Markov decision processes with Wasserstein distance," *IEEE Control Systems Letters*, vol. 1, no. 1, pp. 164–169, 2017.
- [13] L. Huang, J. Walrand, and K. Ramchandran, "Optimal demand response with energy storage management," in *IEEE International Conference on Smart Grid Communications*, 2012.
- [14] J. Qin, Y. Chow, J. Yang, and R. Rajagopal, "Online modified greedy algorithm for storage control under uncertainty," *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 1729–1743, 2016.
- [15] S. Samuelson and I. Yang, "Data-driven distributionally robust control of energy storage systems." [Online]. Available: <http://www-bcf.usc.edu/~inssoonya/DRS.ESS.pdf>
- [16] R. Reemtsen, "Discretization methods for the solution of semi-infinite programming problems," *Journal of Optimization Theory and Applications*, vol. 71, no. 1, pp. 85–103, 1991.
- [17] R. Hettich and K. O. Kortanek, "Semi-infinite programming: Theory, methods, and applications," *SIAM Review*, vol. 35, no. 3, pp. 380–429, 1993.
- [18] M. López and G. Still, "Semi-infinite programming," *European Journal of Operational Research*, vol. 180, pp. 491–518, 2007.