

Optimal Energy Commitments with Storage and Intermittent Supply

Jae Ho Kim

Department of Electrical Engineering, Princeton University, Princeton, New Jersey 08540,
jaek@princeton.edu

Warren B. Powell

Department of Operations Research and Financial Engineering, Princeton University, Princeton, New Jersey 08540,
powell@princeton.edu

We formulate and solve the problem of making advance energy commitments for wind farms in the presence of a storage device with conversion losses, mean-reverting price process, and an autoregressive energy generation process from wind. We derive an optimal commitment policy under the assumption that wind energy is uniformly distributed. Then, the stationary distribution of the storage level corresponding to the optimal policy is obtained, from which the economic value of the storage as the relative increase in the expected revenue due to the existence of storage is obtained.

Subject classifications: Markov decision process; dynamic programming; energy.

Area of review: Stochastic Models.

History: Received September 2009; revisions received May 2010, January 2011; accepted April 2011.

1. Introduction

The emphasis on renewables, such as the goal set by the Department of Energy to have 20% of electric power from wind by 2030, has raised the importance of efficiently managing wind and understanding the factors that affect the cost of using wind. Currently, wind energy accounts for a small fraction in the market, and the grid operators allow the wind energy producers to deliver any amount of energy they produce at a given time. However, as the share of wind energy in the market grows, such a policy will become impractical, and grid operators will need to make commitments on the amount of wind energy that will be delivered in advance. Unfortunately, making commitments is complicated by the inherent uncertainty of wind. This uncertainty can be mitigated by the presence of storage, which also introduces the dimension of losses due to the conversion needed to store and retrieve energy.

We address the problem of making a commitment at time t to deliver energy from wind during the time interval $[t, t + 1)$. The model is most easily applied in the hour-ahead market, although it can be used in an approximate fashion in the day-ahead market. Energy storage has long been recognized as an important technology for smoothing the variability of wind (Castronuovo and Lopez 2004, and Korpaas et al. 2003, García-González et al. 2008, Brunetto and Tina 2007, Ibrahim et al. 2008). We assume that we store energy when the available energy from wind exceeds the commitments we have made, but we may incur significant conversion losses. The problem has cosmetic similarities with classical inventory problems (storing product

to meet demand), but with some fundamental differences. Inventory problems are typically trying to control the supply of product to meet an exogenous demand (Axsäter 2000, Zipkin 2000). In our problem, we have exogenous supply (energy generated from wind) to meet demand by making advance commitments. This problem is similar to the reservoir management problem, which is characterized by random rainfall (see Nandalal and Bogardi 2007 for an excellent review of dynamic programming models for reservoir management). Our problem is distinguished by the need to make advance commitments, along with our interest in a simple, analytical solution that can be used in economic models.

In this paper, we derive an optimal policy for making energy commitments from wind in the presence of an energy storage device. We then use this policy to study the economics of storage capacity in this setting. Given the richness of the application, we analyze a stylized version of the problem, which allows us to derive the optimal policy in a simple, analytic form. Our model captures some important dimensions of the real problem, such as the storage capacity constraints, storage conversion losses, and a mean-reverting process for real-time electricity prices. At the same time, we make a number of simplifying assumptions. Some of these include:

- We assume that we are a small player in a large market, making it possible to sell all of the energy we produce as long as we make advance commitments. In addition, we assume that if the energy from wind (plus what is available

in storage) falls below our commitment, we can make up the entire shortfall using the current spot price.

- If the model is applied in a day-ahead market, we ignore the ability to make hour-ahead adjustments.

- We capture storage capacity and conversion losses, but we otherwise ignore the physics of energy storage, such as the relationship between the rate of storage and storage capacity, and the impact of full discharges on battery life.

- We assume mean-reverting electricity prices and stationarity in the errors in wind forecasts.

- Our analytical model assumes that wind follows a uniform distribution, although we then quantify this error in experimental comparisons using actual wind patterns.

More realistic models require an algorithmic solution. A goal of our research is a simple policy that can be used in economic models without requiring the complex machinery of stochastic optimization algorithms.

The goal of a wind farm operator is to maximize the cumulative profit over time by computing the amount of electricity to commit to sell during the time interval $[t, t + 1)$ at each time t . Brown and Matos (2008), Brunetto and Tina (2007), Castronuovo and Lopez (2004), and Korpaas et al. (2003) attempt to solve the problem by solving a deterministic optimization problem given a particular sample path over a finite horizon and then averaging the results over the sample paths. The sample paths are drawn from a fixed $(T + 1)$ -dimensional distribution describing the electricity generated from the wind farm during the time interval $[t, t + 1)$ for each $t = 0, 1, \dots, T$. However, this approach does not produce a valid, admissible policy. In practice, we need a policy that allows the wind farm operator to compute at time t the amount of electricity to commit to sell during the time interval $[t, t + 1)$ based on the state of the environment at time t . The objective of this paper is to find such a policy and analyze it.

The contributions of this paper are as follows: (1) We derive an analytical expression for an optimal policy, and the value of storage, for a stylized model of an energy storage process in the presence of intermittent generation requiring advance commitments. (2) We establish assumptions on the electricity price and the distribution of wind, size of the storage, and the decision epoch intervals that allow us to derive the optimal policy for energy commitment in a closed form and explain the implications of those assumptions. (3) Under those assumptions, we derive the optimal policy for advance energy commitment in a simple, analytical form, when we have storage with an arbitrary round-trip efficiency, and when electricity prices are mean reverting. The optimal policy obtained under such assumptions resembles the optimal policy for the well-known newsvendor problem (Khouja 1999, Petruzzi and Dada 1999). (4) We obtain the stationary distribution of the storage level corresponding to the optimal policy, from which we find the economic value of the storage as the relative increase in the expected revenue due to the existence

of storage. (5) We test our policy using wind energy generated from truncated Gaussian distributions and demonstrate that the error introduced by assuming a uniform distribution for wind is reasonably small.

This paper is organized as follows. In §2, we model the wind energy storage problem as an MDP with continuous-state and control variables. In §3, we present our assumptions and the structural properties of the optimal value function of the MDP. In §4, the optimal policy for the infinite-horizon problem for a storage with general round-trip efficiency is obtained. Then, the stationary distribution of the storage level corresponding to the optimal policy is obtained, from which the economic value of the storage as the relative increase in revenue due to existence of storage is derived. In §5, we compute the economic value of storage using the wind speed data obtained from the North American Land Data Assimilation System (NLDAS) project (Cosgrove et al. 2003) and the electricity price data provided by a utility company. In §6, we summarize our conclusions.

2. The Model

Operating a wind farm depends on two markets: the electricity spot market and the regulating market. We sell to the spot market and pay a penalty when we fail to meet our commitment. The grid operator buys energy from the regulating market when we fail to meet our commitment. In the spot market, the energy producers make their commitments to deliver (sell) electricity in advance, whereas the regulating market is a marketplace for reserve energy in which the producers have the ability to sell electricity on a shorter notice than the spot market (Korpaas et al. 2003, MacKerron and Pearson 2000, Morthorst 2003). As a wind farm operator, when the electricity production exceeds our expectation and we have an excess amount of electricity left over after fulfilling the contractual commitment, we store the excess amount. On the other hand, when the electricity production falls too short to meet the contractual commitment, we have to pay a premium, a penalty for failing to meet the commitment, while the producers in the regulating market make up for the gap. Therefore, if we commit too much, we can actually lose money. We have revenues from our sale on the spot market and costs from tapping into the regulating market when we fail to meet our commitment for delivery on the spot market (see MacKerron and Pearson 2000, Chapter 16, for a detailed exposition of the market system).

At each time t , the market participants submit their bid for the supply and demand for electricity that must be delivered during the time interval $[t, t + 1)$. The market overseer collects the bidding information and determines the spot market and the regulating market price for the time interval $[t, t + 1)$ shortly after the participants submit their bids. Therefore, as a wind farm operator, we do not know what the prices will be when we are making our commitments.

We make the following assumptions. First, we assume that at each time t we have a probability distribution of the electricity we will generate during the time interval $[t, t + 1)$. Second, we assume that we are a small participant in the market such that the market can always absorb our supply, and the effect of our bidding on the expected spot market and the regulating market prices of the electricity is negligible. Then, the prices can be treated as exogenous variables and we only need to determine the amount of electricity to commit to sell. Third, we assume that the spot market price of the electricity is mean reverting and the ratio of the expected spot market price over the expected regulating market price is always less than the round-trip efficiency of our storage with the discount factor. Otherwise, the cost of using the storage, which can be measured by the conversion loss, will be greater than the expected cost of tapping into the reserve energy in the regulating market, negating the purpose of using a storage device in the first place. The third assumption is crucial in maintaining the concavity of the optimization problem.

2.1. System Parameters

R_{\max} = upper limit on the storage. (*unit: storage energy capacity unit*)

ρ_R = coefficient used to convert the generated electricity to potential energy in the storage. (*unit: storage unit/electricity unit*)

ρ_E = coefficient used to convert the potential energy in the storage to electricity. (*unit: electricity unit/storage unit*)

Note that $0 < \rho_E \rho_R < 1$, where $\rho_E \rho_R$ denotes the round-trip efficiency. Throughout this paper, $1 - \rho_E \rho_R$ is referred to as the conversion loss from storage. $\rho_R \rho_E$ is around 0.6–0.8 for most of the existing storage systems (Sioshanshi et al. 2009).

μ_p = mean of the spot market price of the electricity. (*unit: dollar/electricity unit*)

σ_p = standard deviation of the change in spot market price of the electricity. (*unit: dollar/electricity unit*)

κ = mean-reversion parameter for the spot market price of the electricity. κ is proportional to the expected frequency at which the spot market price crosses the mean per unit time. (*unit: 1/time unit*)

$\Delta\tau$ = time interval between decision epochs.

m = slope of the penalty cost for overcommitment.

b = intercept of the penalty cost for overcommitment. (*unit: dollar/electricity unit*)

That is, when the spot market price of the electricity is p_t , the penalty for overcommitment is $mp_t + b$.

μ_Y = mean of the electricity generated from the wind farm per unit time. (*unit: electricity unit/time unit*)

σ_Y = standard deviation per unit time of the electricity generated from the wind farm. (*unit: electricity unit/time unit*)

γ = discount factor in the MDP model. $0 < \gamma < 1$.

2.2. State Variables

Let $t \in \mathbb{N}_+$ be a discrete time index corresponding to the decision epoch. The actual time corresponding to the time index t is $t\Delta\tau$.

R_t = storage level at time t . $0 \leq R_t \leq R_{\max}$, $\forall t$.

Y_t = electricity generated from the wind turbines during the time interval $[t - 1, t)$. $Y_t \geq 0$, $\forall t$.

p_t = spot market price for electricity delivered during the time interval $[t - 1, t)$. $p_t \geq 0$, $\forall t$.

$W_t = ((Y_{t'})_{1 \leq t' \leq t}, p_t)$ = exogenous state of the system.

$S_t = (R_t, W_t)$ = state of the system at time t .

2.3. Decision (Action) Variable

x_t = amount of electricity we commit to sell on the spot market during the time interval $[t, t + 1)$ determined by signing the contract at time t . $x_t \geq 0$.

Because we are making an advance commitment, x_t is not constrained by R_t . The lack of an upper bound on x_t indicates that we are a small player in the market, and hence there will always be enough demand in the market to absorb our supply as long as we are making an advance commitment.

2.4. Exogenous Process

\hat{y}_t = noise that captures the random evolution of Y_t . Specifically,

$$Y_{t+1} = \mu_Y \Delta\tau + \sum_{i=0}^{M-1} \alpha_i (Y_{t-i} - \mu_Y \Delta\tau) + \hat{y}_{t+1} \quad (1)$$

for some order M and coefficients α_i for $0 \leq i \leq M - 1$. $(\hat{y}_t)_{t \geq 1}$ and $(Y_t)_{t \geq 1}$ must be proportional to $\Delta\tau$.

\hat{p}_t = noise that captures the random evolution of p_t . Specifically, we use a discrete-time version of the Ornstein-Uhlenbeck process:

$$p_{t+1} - p_t = \kappa(\mu_p - p_t)\Delta\tau + \hat{p}_{t+1}.$$

Let Ω be the set of all possible outcomes and let \mathcal{F} be a σ -algebra on the set, with filtrations \mathcal{F}_t generated by the information given up to time t :

$$\mathcal{F}_t = \sigma(S_0, x_0, Y_1, S_1, x_1, Y_2, S_2, x_2, \dots, Y_t, S_t, x_t).$$

\mathbb{P} is the probability measure on the measure space (Ω, \mathcal{F}) . Throughout this paper, a variable with subscript t is unknown (random) at time $t - 1$ and becomes known (deterministic) at time t . In other words, a variable with subscript t is \mathcal{F}_t -measurable. We have defined the state of our system at time t as all variables that are \mathcal{F}_t -measurable and needed to compute our decision at time t .

2.5. Storage Transition Function

$$R_{t+1} = \begin{cases} R_{\max}, & \text{if } R_t + \rho_R(Y_{t+1} - x_t) \geq R_{\max}. \\ R_t + \rho_R(Y_{t+1} - x_t), & \text{if } x_t < Y_{t+1}, R_t + \rho_R(Y_{t+1} - x_t) < R_{\max}. \\ R_t - \frac{1}{\rho_E}(x_t - Y_{t+1}), & \text{if } Y_{t+1} \leq x_t < \rho_E R_t + Y_{t+1}. \\ 0, & \text{if } x_t \geq \rho_E R_t + Y_{t+1}. \end{cases} \quad (2)$$

If Y_{t+1} exceeds the commitment x_t , we store the excess amount $Y_{t+1} - x_t$ with a conversion factor, ρ_R . If Y_{t+1} is less than x_t , the potential energy in the storage must be converted into electricity with a conversion factor, ρ_E , to fulfill the gap, $x_t - Y_{t+1}$. If the amount of electricity generated during the time interval $[t, t + 1]$ plus the electricity that can be obtained by converting the potential energy in storage is not enough to cover the contractual commitment, we deplete our storage, and we have to pay for the gap. It is important to note the difference between the storage transition function shown above and the transition functions that generally appear in traditional inventory management and resource allocation problems (Axsäter 2000, Zipkin 2000). Unlike many of the transition functions that appear in traditional problems, here R_{t+1} is not a concave or convex function of x_t or R_t , which makes the concavity of the optimization problem not obvious.

2.6. Contribution (Revenue) Function

The profit we make during the time interval $[t, t + 1]$ is given by

$$\hat{C}_{t+1} = \begin{cases} p_{t+1}x_t, & \text{if } x_t < \rho_E R_t + Y_{t+1}. \\ p_{t+1}x_t - (mp_{t+1} + b)[x_t - (\rho_E R_t + Y_{t+1})], & \text{if } x_t \geq \rho_E R_t + Y_{t+1}. \end{cases}$$

$p_{t+1}x_t$ is the profit we earn by delivering x_t amount of electricity to the market during the time interval $[t, t + 1]$, and $mp_{t+1} + b$ is the penalty we pay in the case of overcommitment. Assume

$$m \geq \frac{\gamma}{\rho_E \rho_R} \quad \text{and} \quad b \geq \frac{\gamma}{\rho_E \rho_R} \mu_p. \quad (3)$$

Then, the cost of using the storage, which can be measured by the conversion loss, is less than the cost of overcommitment. Otherwise, for the purpose of maximizing the revenue, there will be no reason to use a storage in the first place. This affine penalty is sufficient to ensure the concavity of the stochastic optimization problem. Note that these lower bounds on the penalty factor are unfavorable assumptions—they make the environment in which we

operate more adverse and lead to a more conservative policy. If we have to operate in an environment where the above assumptions do not hold, the optimal policy derived in this paper under the above assumptions may not be optimal in maximizing revenue, but it should still be robust with limited risk—we lose less money than expected in the case of overcommitment. Define

$$\begin{aligned} C(S_t, x_t) &:= \mathbb{E}[\hat{C}_{t+1} | S_t, x_t] \\ &= [\mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p)] \\ &\quad \cdot \left[x_t - m \cdot \int_{0 \leq y \leq x_t - \rho_E R_t} F_t(y) dy \right] \\ &\quad - b \int_{0 \leq y \leq x_t - \rho_E R_t} F_t(y) dy, \end{aligned} \quad (4)$$

where

$$F_t(y) = \mathbb{P}[Y_{t+1} \leq y | \mathcal{F}_t].$$

$C(\cdot)$ is known as the contribution, or the reward function. See §1 in the e-companion to this paper for the derivation of (4). An electronic companion to this paper is available as part of the online version that can be found at <http://orjournal.informs.org/>.

2.7. Objective Function

Let Π be the set of all policies. A policy is an \mathcal{F}_t -measurable function $X^\pi(S_t)$ that describes the mapping from the state at time t , S_t , to the decision at time t , x_t . For each $\pi \in \Pi$, let

$$G_t^\pi(S_t) := \mathbb{E} \left[\sum_{t'=t}^T \gamma^{t'-t} C(S_{t'}, X^\pi(S_{t'})) \mid S_t \right], \quad \forall 0 \leq t \leq T,$$

where $0 < \gamma < 1$ is the discount factor and T indicates the end of the horizon. The objective, then, is to find an optimal policy $\pi = \pi^*$ that satisfies

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

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

3. Main Assumptions and Structural Result

The main contribution of this paper is the closed-form representation of the optimal policy for advance intermittent energy commitments that also allows us to express the value of the energy storage in a closed form. In order to achieve the results, we need assumptions on the probability distribution of the spot market electricity price and wind energy, limit on the storage size, and the decision epoch intervals.

3.1. Electricity Price and Wind Energy

First, we assume that $(\hat{p}_t)_{t \geq 0}$ and $(\hat{y}_t)_{t \geq 1}$ are independent in $(\Omega, \mathcal{F}, \mathbb{P})$. It is well known that the price of the electricity mainly depends on the demand as well as the main source of energy that is controllable, for example, electricity generated from coal plants. It is fairly reasonable to assume that the fluctuation in the electricity price is not significantly influenced by the fluctuation in the uncontrollable and unpredictable energy supply from our wind farm, especially if we are a small player in the market. In most cases, intermittent energy plays a minor role in the electricity markets, anyway.

Next, assume $(\hat{p}_t)_{t \geq 0}$ are i.i.d. with distribution $\mathcal{N}(0, \sigma_p^2)$. Then, $(p_t)_{t \geq 0}$ is a standard mean-reverting process and

$$\mathbb{E}[p_{t+n} | \mathcal{F}_t] = \mu_p + (1 - \kappa \Delta \tau)^n (p_t - \mu_p), \quad \forall n, t \in \mathbb{N}_+. \quad (5)$$

It is common to use a mean-reverting process to model electricity prices, as shown in Eydeland and Wolyniec (2003). Similarly, assume $(\hat{y}_t)_{t \geq 1}$ are 0-mean and i.i.d. with standard deviation $\sigma_Y \Delta \tau$. Then, in most cases, the distributions of $(\hat{y}_t)_{t \geq 1}$ are assumed to be truncated Gaussian with mean 0 and standard deviation $\sigma_Y \Delta \tau$. However, in this paper we assume that $(\hat{y}_t)_{t \geq 1}$ are uniformly distributed with mean 0 with standard deviation $\sigma_Y \Delta \tau$. Assuming that $(\hat{y}_t)_{t \geq 1}$ are uniformly distributed allows us to explicitly compute various expectations that are needed to derive the optimal policy in a closed form. Because a truncated Gaussian distribution is bounded, as long as we match the mean and the variance, a uniform distribution can be a statistically robust substitute for the truncated Gaussian distribution in the context of optimizing a value function. This fact is demonstrated in §5 where we conduct numerical experiments in which we apply the optimal policy derived under the assumption of uniformly distributed $(\hat{y}_t)_{t \geq 1}$ to the data generated from truncated Gaussian distributions. Then, given \mathcal{F}_t , $Y_{t+1} \sim \mathcal{U}(\theta_t, \theta_t + \beta)$, where

$$\beta := 2\sqrt{3}\sigma_Y \Delta \tau$$

and

$$\theta_t := \mu_Y \Delta \tau + \sum_{i=0}^{M-1} \alpha_i (Y_{t-i} - \mu_Y \Delta \tau) - \frac{\beta}{2}, \quad \forall t. \quad (6)$$

The cumulative density function (CDF) of Y_{t+1} computed at time t is given by

$$F_t(y) = \mathbb{P}[Y_{t+1} \leq y | \mathcal{F}_t] = \begin{cases} 0, & \text{if } y < \theta_t, \\ \frac{y - \theta_t}{\beta} & \text{if } \theta_t \leq y \leq \theta_t + \beta \\ 1, & \text{if } y > \theta_t + \beta. \end{cases}$$

The expected contribution function $C(\cdot)$ is not indexed by t because the CDF $F_t(\cdot)$ is determined by θ_t , which is a deterministic function of S_t . The expected contribution is completely determined by S_t and x_t . However, it is important to note that θ_t and β do not necessarily have

to be defined as shown above. The results obtained in this paper are applicable as long as we use a forecasting model that predicts that the electricity produced during the time interval $[t, t + 1)$ is uniformly distributed given \mathcal{F}_t .

3.2. Size of the Storage

Next, we need an assumption on the size of the storage. Because the electricity price is mean reverting, if we have an infinitely large storage, a naive policy that stores the energy when the expected spot market price is less than some fixed price and commits to sell the energy in storage plus the energy we are certain to produce when the spot market price is greater than some fixed price will be a riskless arbitrage policy. Arbitrage here means that there is zero probability of losing money due to overcommitment or losing energy due to the storage being full. There is always a significant conversion loss. Such a case is comparable to trading a stock whose price is mean reverting. In reality, a storage with reasonably good round-trip efficiency that can be charged and discharged in a short amount of time will be expensive to build and maintain, and we need an intelligent way of determining the appropriate size of the storage. We propose that the size of the storage be determined in comparison to ν , given by:

$$\nu := \rho_R \frac{\sigma_Y}{\kappa} 2\sqrt{3} \min \left[\frac{m-1}{m}, \frac{b}{b + \rho_E \rho_R \gamma \mu_p} \right]. \quad (7)$$

It is obvious that as the penalty factors m and b become larger, we need to allow for a larger storage because our commitment level will be more conservative and we will end up storing more energy. Also, if the round-trip efficiency of the storage $\rho_E \rho_R$ is small, we must allow for a larger storage in order to compensate for the energy that will be lost in conversion. Next, because $\gamma \mu_p$ is the discounted expected spot market price of the electricity, if $\gamma \mu_p$ is small, we need to allow for a larger storage because our commitment level will be more conservative.

What makes ν interesting is the term σ_Y / κ . Recall that κ is proportional to the expected number of times the price crosses the mean per-unit time. Then, $1/\kappa$ is proportional to the expected amount of time between two consecutive crossings. Therefore, σ_Y / κ is proportional to the volatility in the wind energy that is produced while the spot market price “completes a cycle.” Because R_{\max} determines our ability to accumulate energy while the price moves, we must allow for a larger storage when σ_Y / κ gets larger. If $R_{\max} = \infty$, we can implement an arbitrage policy, as explained above. If $R_{\max} \leq \nu$, we must implement a more active risk taking policy that considers the movement of the price towards the mean but not “count on” the price reaching a desirable level within a desirable amount of time. The middle regime in which $\nu < R_{\max} < \infty$ will demand the most complicated policy that mixes risk-taking with arbitrage. Finding the optimal policy in this middle regime

will be an interesting research topic, but it is beyond the scope of this paper. For this paper, we assume

$$R_{\max} \leq \rho_R \frac{\sigma_Y}{\kappa} 2\sqrt{3} \min \left[\frac{m-1}{m}, \frac{b}{b + \rho_E \rho_R \gamma \mu_p} \right]. \quad (8)$$

(8) is necessary in order to derive (10) shown in the next section, which in turn is necessary to prove lemma (18) that is used to derive the marginal value function in a closed form. However, even though (8) is imposed for mathematical convenience, numbers come out reasonable, as shown in §5. If we use real data to obtain $\sigma_Y, \kappa, \mu_p, \rho_E \rho_R$ and use m and b that satisfy (3), if we let

$$\frac{R_{\max}}{\sigma_Y \rho_R} = 1.7,$$

for example, (8) is satisfied. That is, we can have the size of R_{\max} in the same order of magnitude of the standard deviation in wind energy. Having a storage of limited size allows us to obtain the optimal policy in a closed form and provide us with various insights, as is shown in §4. Moreover, before investing a significant amount of capital to build a large storage, it is reasonable to assume that wind farm operators will start with a small storage, study its effects, and then subsequently make the investment for additional storage. This paper derives the optimal policy for energy commitment and the corresponding value of the storage when the storage is small.

As will be shown in §4, the optimal policy under the assumption (8) will still depend on the mean of the electricity price and how far the price is away from the mean. However, the optimal policy will be based on the premise that the storage is not large enough to allow us to avoid the risk of overcommitment by waiting for the price to rise without facing the risk of losing energy due to the storage being full. Thus, (8) forces us to always balance the risk of overcommitment and the risk of undercommitment. We not only want to avoid paying the penalty for overcommitment, but we also want to avoid committing too little and losing energy due to conversion and the storage being full.

Suppose we have a large storage device and (8) is violated, but we choose to implement the policy derived in this paper that is optimal under the assumption of small storage, anyway. Then, the cost of overcommitment will not change, but the risk of undercommitment will be smaller than expected because we are less likely to lose energy due to the storage being full. Therefore, the optimal policy derived in this paper will still be robust when the assumption (8) does not hold.

3.3. Decision Epoch Interval

Finally, we need an assumption on how often we make our commitment decisions. We can rearrange the terms from (8) to obtain

$$\max \left[\frac{R_{\max}(m - \rho_E \rho_R \gamma)}{2\sqrt{3}(m-1)\rho_R \sigma_Y - R_{\max} \rho_E \rho_R \gamma \kappa}, \frac{R_{\max} b}{2\sqrt{3}\rho_R \sigma_Y b - R_{\max} \rho_E \rho_R \gamma \kappa \mu_p} \right] \leq \frac{1}{\kappa}.$$

We assume that the time interval $\Delta\tau$ between our decision epochs satisfies the following:

$$\max \left[\frac{R_{\max}(m - \rho_E \rho_R \gamma)}{2\sqrt{3}(m-1)\rho_R \sigma_Y - R_{\max} \rho_E \rho_R \gamma \kappa}, \frac{R_{\max} b}{2\sqrt{3}\rho_R \sigma_Y b - R_{\max} \rho_E \rho_R \gamma \kappa \mu_p} \right] \leq \Delta\tau \leq \frac{1}{\kappa}. \quad (9)$$

(9) ensures that the price always moves toward the mean in expectation, but does not overshoot and move pass the mean in expectation. The lower bound can be rearranged to be written as

$$R_{\max} \leq \rho_R \beta \min \left[\frac{m-1}{m - \rho_E \rho_R \gamma (1 - \kappa \Delta\tau)}, \frac{b}{b + \rho_E \rho_R \gamma \kappa \Delta\tau \mu_p} \right]. \quad (10)$$

Because we have a limit on the size of our storage, as shown in our assumption (8), if $\Delta\tau$ is too large, the amount of electricity that is produced between our decisions can be too large and we are likely to lose energy due to the storage being full. (9) gives us a reasonable decision epoch time interval $\Delta\tau$.

3.4. Structural Results

In this section, we show some structural results of the value function. Let $V_t^\pi(S_t)$ be a function that satisfies

$$\begin{aligned} V_T^\pi(S_T) &= C(S_T, X_T^\pi(S_T)), \\ V_t^\pi(S_t) &= C(S_t, X_t^\pi(S_t)) + \gamma \mathbb{E}[V_{t+1}^\pi(S_{t+1}) | S_t], \\ &\quad \forall 0 \leq t \leq T-1. \end{aligned}$$

Then, $V_t^\pi(S_t) = G_t^\pi(S_t), \forall 0 \leq t \leq T$. For $0 \leq t \leq T$, let $V_t(S_t)$ satisfy the following:

$$\begin{aligned} V_T(S_T) &= \max_{x \in \mathbb{R}_+} C(S_T, x), \\ V_t(S_t) &= \max_{x \in \mathbb{R}_+} \{C(S_t, x) + \gamma \mathbb{E}[V_{t+1}(S_{t+1}) | S_t, x]\}, \\ &\quad \forall 0 \leq t \leq T-1. \end{aligned}$$

$V_t(S_t)$ is known as the *value function*. According to Puterman (1994), $V_t(S_t) = G_t^{\pi^*}(S_t), \forall 0 \leq t \leq T$. Denote

$$V_t^x(S_t, x) := \mathbb{E}[V_{t+1}(S_{t+1}) | S_t, x], \quad \forall 0 \leq t \leq T.$$

The augmented value function $V_t^x(S_t, x)$ is an example of a *Q-factor*. Let

$$\begin{aligned} x_t^* &:= \arg \max_{x \in \mathbb{R}_+} \{C(S_t, x) + \gamma V_t^x(S_t, x)\} \\ &= X^{\pi^*}(S_t), \quad \forall 0 \leq t \leq T. \end{aligned}$$

Then,

$$\begin{aligned} V_t(S_t) &= \max_{x \in \mathbb{R}_+} \{C(S_t, x) + \gamma V_t^x(S_t, x)\} \\ &= C(S_t, x_t^*) + \gamma V_t^x(S_t, x_t^*), \quad \forall 0 \leq t \leq T. \end{aligned}$$

At the end of the horizon, we can show that

$$\frac{d}{dR_T} V_T(S_T) = \rho_E [\mu_p + (1 - \kappa \Delta \tau)(p_T - \mu_p)], \quad (11)$$

and hence

$$\frac{d^2}{dR_T^2} V_T(S_T) = 0. \quad (12)$$

See §2 in the e-companion to this paper for the derivation of (11) and (12).

Now that we have defined the value function, we present its structure. The structural results are mainly attributable to the storage transition function and the contribution function, and they follow from three of the aforementioned assumptions: $(\hat{y}_t)_{t \geq 1}$ and $(\hat{p}_t)_{t \geq 1}$ are independent, $(p_t)_{t \geq 1}$ is mean reverting as shown in (5), and

$$m \geq \frac{\gamma}{\rho_E \rho_R} \quad \text{and} \quad b \geq \frac{\gamma}{\rho_E \rho_R} \mu_p.$$

Then, $\forall 0 \leq t \leq T - 1$, we have:

STRUCTURAL RESULT 1. $C(S_t, x) + \gamma V_t^x(S_t, x)$ is a concave function of (R_t, x) .

STRUCTURAL RESULT 2. The optimal decision x_t^* is positive and finite and

$$\frac{\partial}{\partial x} C(S_t, x_t^*) + \gamma \frac{\partial}{\partial x} V_t^x(S_t, x_t^*) = 0. \quad (13)$$

STRUCTURAL RESULT 3.

$$\begin{aligned} \frac{d}{dR_t} V_t(S_t) &= \rho_E [\mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p)] \\ &+ \gamma \mathbb{E} \left[\left(\rho_E \frac{\partial R_{t+1}}{\partial x} + \frac{\partial R_{t+1}}{\partial R_t} \right) \frac{d}{dR_{t+1}} V_{t+1}(S_{t+1}) \Big| S_t, x_t^* \right]. \end{aligned} \quad (14)$$

STRUCTURAL RESULT 4. $V_t(S_t)$ is a concave function of R_t .

STRUCTURAL RESULT 5.

$$\begin{aligned} \rho_E [\mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p)] \\ \leq \frac{d}{dR_t} V_t(S_t) \leq \frac{1}{\rho_R} [\mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p)]. \end{aligned} \quad (15)$$

See §2 in the e-companion to this paper for the proof of the above results.

In Structural Result 3, which shows the recursive relationship between the marginal value functions, the meaning of the term

$$\rho_E [\mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p)]$$

is obvious; if we had an extra ΔR_t amount of energy in storage, we can commit to sell it and gain

$$\Delta R_t \rho_E [\mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p)]$$

in expected revenue. However, the second term requires some analysis. From (2), we know that

$$\begin{aligned} \rho_E \frac{\partial R_{t+1}}{\partial x} + \frac{\partial R_{t+1}}{\partial R_t} \Big|_{x=x_t^*} \\ = \begin{cases} 1 - \rho_E \rho_R, & \text{if } x_t^* < Y_{t+1}, R_t + \rho_R(Y_{t+1} - x_t^*) < R_{\max}, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

describes the conversion loss that occurs when we use the energy that is put into the storage when we generate more electricity than we need to satisfy the commitment. Therefore, the term

$$\mathbb{E} \left[\left(\rho_E \frac{\partial R_{t+1}}{\partial x} + \frac{\partial R_{t+1}}{\partial R_t} \right) \frac{d}{dR_{t+1}} V_{t+1}(S_{t+1}) \Big| S_t, x_t^* \right]$$

can be seen as the expected portion of the marginal future value function that is saved by not having to go through the process of energy conversion.

4. Main Result—Infinite-Horizon Analysis

In this section, we derive the marginal value function and the corresponding optimal policy for advance energy commitment that maximizes the expected revenue in the infinite-horizon case. However, although we can obtain the value of always having a storage as shown in this paper, it is important to note that the cost of always having storage is not the cost of installing the storage once in the beginning. Batteries have finite lifetime, and we might have to reinstall them every 10 years, for example. We let $T \rightarrow \infty$ and drop the index t from the value function:

$$V(S_t) = \lim_{T \rightarrow \infty} \mathbb{E} \left[\sum_{t'=t}^T \gamma^{t'-t} C(S_{t'}, X^{\pi^*}(S_{t'})) \Big| S_t \right].$$

Then, $V(S_t)$ satisfies

$$\begin{aligned} V(S_t) &= \max_{x \in \mathbb{R}_+} \{C(S_t, x) + \gamma \mathbb{E}[V(S_{t+1}) \Big| S_t, x]\} \\ &= C(S_t, x_t^*) + \gamma V^x(S_t, x_t^*). \end{aligned}$$

Because the structural properties shown in the previous section hold true for all T , $V(S_t)$ maintains those structural properties. In §4.1, we derive the optimal policy using the

main assumptions stated in §2 and the structural results shown in §3. We first state

THEOREM 1. *The optimal policy, when the electricity generated from the wind farm is uniformly distributed from θ_t to $\theta_t + \beta$, is given by*

$$x_t^* = X^{\pi^*}(S_t) \\ = \rho_E R_t + \theta_t + \frac{\mu_p K_1 + (p_t - \mu_p)(1 - \kappa \Delta \tau) K_2}{m[\mu_p + (p_t - \mu_p)(1 - \kappa \Delta \tau)] + b} \beta \quad (16)$$

where

$$K_1 = 1 - \gamma \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \left(\exp \left[\gamma (1 - \rho_R \rho_E) \frac{1}{\beta} \frac{R_{\max}}{\rho_R} \right] - 1 \right),$$

and

$$K_2 = 1 - \gamma (1 - \kappa \Delta \tau) \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \\ \cdot \left(\exp \left[\gamma (1 - \kappa \Delta \tau) (1 - \rho_R \rho_E) \frac{1}{\beta} \frac{R_{\max}}{\rho_R} \right] - 1 \right).$$

Before proving (16), we first analyze its components. Because $\rho_E R_t$ is the amount of electricity that can be produced by converting the energy in storage and θ_t is the amount of electricity that is certain to be produced, $\rho_E R_t + \theta_t$ can be seen as the riskless term. Because there is a limit on the size of the storage and we lose energy if the storage is full, we always want to commit to sell at least $\rho_E R_t + \theta_t$. The issue is then how much more to commit relative to this base level. Overcommitment is costly because the expected penalty always exceeds the expected spot price. Undercommitment is costly for two reasons. First, excess production must be stored and storage is not free because the round-trip efficiency is less than 1. Second, because there is a limit on the amount of energy you can store, R_{\max} , if we commit too little and produce too much we lose the production that cannot be stored. Therefore, the optimal extra commitment over the base level must balance the cost of overcommitment and undercommitment. β is the uncertainty in the electricity production, and committing

$$\frac{\mu_p K_1 + (p_t - \mu_p)(1 - \kappa \Delta \tau) K_2}{m[\mu_p + (p_t - \mu_p)(1 - \kappa \Delta \tau)] + b} \quad (17)$$

fraction of β achieves the balance between the cost of overcommitment and the cost of undercommitment. Note that the solution to a typical newsvendor problem states that the vendor should always try to satisfy a fixed fraction of the random demand (Khouja 1999, Petruzzi and Dada 1999). However, in our case, the fraction is a function of the price because we can speculate on the movement of the price that is mean reverting.

In §4.2, we obtain the stationary distribution of the storage level corresponding to the optimal policy. In §4.3, we derive the economic value of the storage as the relative increase in average revenue due to the existence of the storage.

4.1. Optimal Policy

In this section, we prove the optimal policy (16) by first deriving the marginal value function. From Structural Result 2, we know that the optimal decision x_t^* must satisfy

$$\frac{\partial}{\partial x} C(S_t, x_t^*) + \frac{\partial}{\partial x} V^x(S_t, x_t^*) \\ = \frac{\partial}{\partial x} C(S_t, x_t^*) + \gamma \mathbb{E} \left[\frac{\partial R_{t+1}}{\partial x} \frac{d}{dR_{t+1}} V(S_{t+1}) \mid S_t, x_t^* \right] \\ = 0.$$

Therefore, in order to compute x_t^* , we only need to know the derivative of $V(S_{t+1})$ with respect to R_{t+1} , and we do not need to know $V(S_{t+1})$ itself. To derive $(d/dR_{t+1})V(S_{t+1})$, we need the following lemma:

LEMMA 1.

$$x_t^* + \frac{R_{\max} - R_t}{\rho_R} \leq \theta_t + \beta, \quad \forall t. \quad (18)$$

PROOF. See §3 in the e-companion to this paper. The proof utilizes the inequality (10). \square

We know that $\theta_t + \beta - x_t^*$ is the maximum amount of excess electricity that can be left over after fulfilling the commitment. Suppose that the inequality (18) does not hold. Then, $\rho_R(\theta_t + \beta - x_t^*) \leq R_{\max} - R_t$, indicating that there is always enough room left in the storage to accommodate all of the excess electricity, implying that there is no risk of undercommitment at all. However, we have restricted the size of the storage as shown in (10) precisely to avoid such a situation. We know that the optimal policy ought to balance the risk of undercommitment and the risk of overcommitment. The above lemma allows us to compute $(d/dR_t)V(S_t)$, from which we can derive the optimal policy. We first state

THEOREM 2.

$$\frac{d}{dR_t} V(S_t) = \rho_E \mu_p \exp \left[\gamma (1 - \rho_R \rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_t}{\rho_R} \right) \right] \\ + \rho_E (p_t - \mu_p) (1 - \kappa \Delta \tau) \\ \cdot \exp \left[\gamma (1 - \kappa \Delta t) (1 - \rho_E \rho_R) \frac{1}{\beta} \left(\frac{R_{\max} - R_t}{\rho_R} \right) \right]. \quad (19)$$

PROOF. Here we show a condensed version of the proof by omitting various algebraic steps. See §4 in the e-companion to this paper for a detailed proof. We prove the theorem by using backward induction in the finite-horizon setting and letting T go to infinity. First, we make the induction hypothesis that

$$\frac{d}{dR_{T-i}} V_{T-i}(S_{T-i}) \\ = \rho_E \mu_p \sum_{j=0}^i \frac{1}{j!} \left[\gamma (1 - \rho_R \rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-i}}{\rho_R} \right) \right]^j$$

$$+ \rho_E(p_{T-i} - \mu_p)(1 - \kappa\Delta\tau) \cdot \sum_{j=0}^i \frac{1}{j!} \left[\gamma(1 - \kappa\Delta t)(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-i}}{\rho_R} \right) \right] \quad (20)$$

for some $i \geq 0$, and prove that

$$\begin{aligned} & \frac{d}{dR_{T-(i+1)}} V_{T-(i+1)}(S_{T-(i+1)}) \\ &= \rho_E \mu_p \sum_{j=0}^{i+1} \frac{1}{j!} \left[\gamma(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-i}}{\rho_R} \right) \right]^j \\ & \quad + \rho_E(p_{T-(i+1)} - \mu_p)(1 - \kappa\Delta\tau) \\ & \quad \cdot \sum_{j=0}^{i+1} \frac{1}{j!} \left[\gamma(1 - \kappa\Delta t)(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-(i+1)}}{\rho_R} \right) \right]^j. \end{aligned}$$

From (11), we know that

$$\frac{d}{dR_T} V_T(S_T) = \rho_E \mu_p + \rho_E(p_T - \mu_p)(1 - \kappa\Delta\tau).$$

Therefore, the expression for $(d/dR_{T-i})V_{T-i}(S_{T-i})$ shown above is true for $i = 0$. From (2), we can show that

$$\begin{aligned} & \left(\rho_E \frac{\partial R_{t+1}}{\partial x} + \frac{\partial R_{t+1}}{\partial R_t} \right) \frac{1}{j!} \left(\frac{R_{\max} - R_{t+1}}{\rho_R} \right)^j \Big|_{x=x_t^*} \\ &= \begin{cases} (1 - \rho_E\rho_R) \frac{1}{j!} \left[\frac{R_{\max} - R_t}{\rho_R} - (Y_{t+1} - x_t^*) \right]^j, \\ \quad \text{if } x_t^* < Y_{t+1}, R_t + \rho_R(Y_{t+1} - x_t^*) < R_{\max} \\ 0, \quad \text{otherwise.} \end{cases} \end{aligned}$$

Next, by (18),

$$f_i(y) = \frac{1}{\beta}, \quad \forall x_t^* \leq y \leq x_t^* + \frac{R_{\max} - R_t}{\rho_R}.$$

Then, we can show

$$\begin{aligned} & \mathbb{E} \left[\left(\rho_E \frac{\partial R_{t+1}}{\partial x} + \frac{\partial R_{t+1}}{\partial R_t} \right) \frac{1}{j!} \left(\frac{R_{\max} - R_{t+1}}{\rho_R} \right)^j \Big| S_t, x_t^* \right] \\ &= (1 - \rho_E\rho_R) \frac{1}{\beta} \frac{1}{(j+1)!} \left(\frac{R_{\max} - R_t}{\rho_R} \right)^{j+1}. \end{aligned}$$

From Structural Result 3,

$$\begin{aligned} & \frac{d}{dR_{T-(i+1)}} V_{T-(i+1)}(S_{T-(i+1)}) \\ &= \rho_E(\mu_p + (1 - \kappa\Delta\tau)(p_{T-(i+1)} - \mu_p)) \\ & \quad + \gamma \mathbb{E} \left[\left(\rho_E \frac{\partial R_{T-i}}{\partial x} + \frac{\partial R_{T-i}}{\partial R_{T-(i+1)}} \right) \cdot \frac{d}{dR_{T-i}} V_{T-i}(S_{T-i}) \Big| S_{T-(i+1)}, x_{T-(i+1)}^* \right] \\ &= \rho_E \mu_p \sum_{j=0}^{i+1} \frac{1}{j!} \left[\gamma(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-(i+1)}}{\rho_R} \right) \right]^j \\ & \quad + \rho_E(p_{T-(i+1)} - \mu_p)(1 - \kappa\Delta\tau) \\ & \quad \cdot \sum_{j=0}^{i+1} \frac{1}{j!} \left[\gamma(1 - \kappa\Delta\tau) \cdot (1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-(i+1)}}{\rho_R} \right) \right]^j. \end{aligned}$$

Therefore, (20) is true for $\forall i \geq 0$. Next, substitute t for $T - (i + 1)$. Then,

$$\begin{aligned} & \frac{d}{dR_t} V_t(S_t) \\ &= \rho_E \mu_p \sum_{j=0}^{T-t} \frac{1}{j!} \left[\gamma(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_t}{\rho_R} \right) \right]^j \\ & \quad + \rho_E(p_t - \mu_p)(1 - \kappa\Delta\tau) \\ & \quad \cdot \sum_{j=0}^{T-t} \frac{1}{j!} \left[\gamma(1 - \kappa\Delta\tau) \cdot (1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_{T-(i+1)}}{\rho_R} \right) \right]^j, \end{aligned}$$

$\forall t \leq T$. If we let T go to infinity,

$$\begin{aligned} & \frac{d}{dR_t} V(S_t) \\ &= \lim_{T \rightarrow \infty} \frac{d}{dR_t} V_t(S_t) \\ &= \rho_E \mu_p \exp \left[\gamma(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_t}{\rho_R} \right) \right] \\ & \quad + \rho_E(p_t - \mu_p)(1 - \kappa\Delta\tau) \\ & \quad \cdot \exp \left[\gamma(1 - \kappa\Delta\tau)(1 - \rho_R\rho_E) \frac{1}{\beta} \left(\frac{R_{\max} - R_t}{\rho_R} \right) \right], \quad \forall t. \quad \square \end{aligned}$$

To compute the optimal decision x_t^* at time t , all we need to know is $(d/dR_{t+1})V(S_{t+1})$. Because we now know what $(d/dR_{t+1})V(S_{t+1})$ is, we are ready to prove (16).

PROOF of (16). From (19),

$$\begin{aligned} & \mathbb{E} \left[\frac{\partial R_{t+1}}{\partial x} \frac{d}{dR_{t+1}} V(S_{t+1}) \Big| S_t, x \right] \\ &= -\mu_p \frac{\rho_R\rho_E}{1 - \rho_R\rho_E} \left(\exp \left[\gamma(1 - \rho_R\rho_E) \frac{1}{\beta} \frac{R_{\max}}{\rho_R} \right] - 1 \right) \\ & \quad - (p_t - \mu_p)(1 - \kappa\Delta\tau)^2 \frac{\rho_R\rho_E}{1 - \rho_R\rho_E} \\ & \quad \cdot \left(\exp \left[\gamma(1 - \kappa\Delta\tau) \cdot (1 - \rho_R\rho_E) \frac{1}{\beta} \frac{R_{\max}}{\rho_R} \right] - 1 \right), \quad (21) \end{aligned}$$

$\forall x \geq \rho_E R_t + \theta_t$. To see the derivation of (21), see §5 in the e-companion to this paper. Then, from Structural Result 2, we know that the optimal decision x_t^* must satisfy

$$\begin{aligned} & \frac{\partial}{\partial x} C(S_t, x_t^*) + \gamma \frac{\partial}{\partial x} V^x(S_t, x_t^*) \\ &= p_{t,t+1} - (mp_{t,t+1} + b)F_t(x_t^* - \rho_E R_t) \\ & \quad + \gamma \mathbb{E} \left[\frac{\partial R_{t+1}}{\partial x} \frac{d}{dR_{t+1}} V(S_{t+1}) \Big| S_t, x_t^* \right] \\ &= 0, \end{aligned}$$

where

$$p_{t,t+1} := \mathbb{E}[p_{t+1} | \mathcal{F}_t] = \mu_p + (1 - \kappa\Delta\tau)(p_t - \mu_p).$$

Therefore,

$$\begin{aligned} & (mp_{t,t+1} + b) \frac{1}{\beta} (x_t^* - \rho_E R_t - \theta_t) \\ &= p_{t,t+1} + \gamma \mathbb{E} \left[\frac{\partial R_{t+1}}{\partial x} \frac{d}{dR_{t+1}} V(S_{t+1}) \mid S_t, x_t^* \right] \\ &= \mu_p + (p_t - \mu_p)(1 - \kappa \Delta \tau) \\ &\quad - \gamma \mu_p \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \left(\exp \left[\gamma(1 - \rho_R \rho_E) \frac{1}{\beta} \frac{R_{\max}}{\rho_R} \right] - 1 \right) \\ &\quad - \gamma (p_t - \mu_p)(1 - \kappa \Delta \tau)^2 \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \\ &\quad \cdot \left(\exp \left[\gamma(1 - \kappa \Delta \tau)(1 - \rho_R \rho_E) \frac{1}{\beta} \frac{R_{\max}}{\rho_R} \right] - 1 \right) \\ &= \mu_p K_1 + (p_t - \mu_p)(1 - \kappa \Delta \tau) K_2. \end{aligned}$$

Then,

$$x_t^* = \rho_E R_t + \theta_t + \frac{\mu_p K_1 + (p_t - \mu_p)(1 - \kappa \Delta \tau) K_2}{m[\mu_p + (p_t - \mu_p)(1 - \kappa \Delta \tau)] + b} \beta. \quad \square$$

Note that both K_1 and K_2 increases when $\rho_R \rho_E$, R_{\max} , or γ is reduced. We know that the optimal amount should naturally depend on the penalty, the round-trip efficiency, the maximum storage limit, and the discount factor as follows. First, it should decrease with increasing penalty, because being short incurs the penalty. Second, it should increase with reduced round-trip efficiency, because being long implies paying to store (losing energy). Third, it should decrease with increasing maximum storage. If our storage capacity is greater, we lose less of the energy we do not sell, and we can afford to be more conservative and commit less. Fourth, it should increase with decreasing discount factor, because the value of what we store now to use in the future decreases with the discount factor.

Next, suppose storage devices with sufficient capacities become ubiquitous in the future, and hence electricity becomes a very liquid asset just like stocks. Then, arbitraguers taking advantage of predictable patterns such as mean reversion will cause the electricity prices to behave more and more like a martingale.

COROLLARY 1. *If $\kappa = 0$, implying that $(p_t)_{t \geq 0}$ is a martingale, then $K_1 = K_2$ and*

$$x_t^* = \rho_E R_t + \theta_t + \frac{p_t}{mp_t + b} K_1 \beta. \quad (22)$$

If the price is a martingale, it is “stochastically constant” and we cannot speculate on the future movement of the price. Then, the fraction is just directly proportional to the ratio between the current expected spot market price and the penalty price.

4.2. Stationary Distribution of the Storage Level

Now that we have the optimal policy (16), we want to assess the expected value of storage corresponding to the

policy. In order to obtain a closed-form expression for the expected value of storage, we must analyze the dynamics of our system at the steady state and derive the stationary distribution of the storage level. Denote

$$Z_t := \frac{\mu_p K_1 + (p_t - \mu_p)(1 - \kappa \Delta \tau) K_2}{m[\mu_p + (p_t - \mu_p)(1 - \kappa \Delta \tau)] + b}, \quad \forall t.$$

From (2), we know that R_{t+1} is a function of (R_t, x_t^*, Y_{t+1}) . Because Y_{t+1} is a function of θ_t and x_t^* is a function of (R_t, Z_t, θ_t) as shown in (16), we can think of R_{t+1} as a function of (R_t, Z_t, θ_t) . However, because θ_t is the amount of electricity that we are certain to produce and commit, we know that R_{t+1} in fact does not depend on θ_t . Thus, R_{t+1} is a function of (R_t, Z_t) . Therefore, if the random process $(Z_t)_{t \geq 0}$ is stationary ergodic, the process $(R_t)_{t \geq 0}$ will reach a steady state. Because $(Z_t)_{t \geq 1}$ is driven by $(p_t)_{t \geq 1}$, we first need to know the distribution of $(p_t)_{t \geq 1}$ in steady state. Here we use the term “steady state” to refer to the unconditional process.

PROPOSITION 1. *At steady state,*

$$p_t \sim \mathcal{N} \left(\mu_p, \frac{\sigma_p^2}{1 - (1 - \kappa \Delta \tau)^2} \right). \quad (23)$$

PROOF. We know that (23) is true if and only if (23) implies

$$p_{t+1} \sim \mathcal{N} \left(\mu_p, \frac{\sigma_p^2}{1 - (1 - \kappa \Delta \tau)^2} \right).$$

Suppose (23) is true. Then,

$$(1 - \kappa \Delta \tau)(p_t - \mu_p) \sim \mathcal{N} \left(0, \frac{(1 - \kappa \Delta \tau)^2 \sigma_p^2}{1 - (1 - \kappa \Delta \tau)^2} \right).$$

Because \hat{p}_{t+1} is independent from p_t and $\hat{p}_{t+1} \sim \mathcal{N}(0, \sigma_p^2)$,

$$(1 - \kappa \Delta \tau)(p_t - \mu_p) + \hat{p}_{t+1} \sim \mathcal{N} \left(0, \frac{\sigma_p^2}{1 - (1 - \kappa \Delta \tau)^2} \right).$$

Then,

$$\begin{aligned} p_{t+1} &= \mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p) \\ &\quad + \hat{p}_{t+1} \sim \mathcal{N} \left(\mu_p, \frac{\sigma_p^2}{1 - (1 - \kappa \Delta \tau)^2} \right). \quad \square \end{aligned}$$

Because Z_t is a deterministic function of p_t , $(Z_t)_{t \geq 1}$ reaches steady state when $(p_t)_{t \geq 1}$ reaches steady state. Although real-time electricity spot prices can be negative due to tax subsidies, it will be extremely rare for the day-ahead forecast price to be negative. Thus, in practice, we may assume that the price is always going to be positive. The first and second moments of Z_t at steady state given $p_t \geq 0$ is

$$\bar{Z}_1 := \mathbb{E} \left[\frac{\mu_p K_1 + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau) K_2}{m[\mu_p + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau)] + b} \mid \varepsilon \geq 0 \right], \quad (24)$$

and

$$\bar{Z}_2 := \mathbb{E} \left[\left(\frac{\mu_p K_1 + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau) K_2}{m[\mu_p + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau)] + b} \right)^2 \middle| \varepsilon \geq 0 \right] \quad (25)$$

where

$$\varepsilon \sim \mathcal{N} \left(\mu_p, \frac{\sigma_p^2}{1 - (1 - \kappa \Delta \tau)^2} \right).$$

Also, define

$$\bar{Z}_1 := \mathbb{E} \left[\frac{\mu_p + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau)}{m[\mu_p + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau)] + b} \middle| \varepsilon \geq 0 \right],$$

and

$$\tilde{Z}_2 := \mathbb{E} \left[\left(\frac{\mu_p + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau)}{m[\mu_p + (\varepsilon - \mu_p)(1 - \kappa \Delta \tau)] + b} \right)^2 \middle| \varepsilon \geq 0 \right],$$

corresponding to the case where $R_{\max} = 0$, which makes $K_1 = K_2 = 1$. $\bar{Z}_1, \bar{Z}_2, \tilde{Z}_1$, and \tilde{Z}_2 can be easily computed via Monte Carlo simulation using sample realizations of ε greater than zero.

PROPOSITION 2. *Then, the stationary distribution of R_t corresponding to the steady state is*

$$\begin{aligned} f_{R_t}(r) &= \frac{d}{dr} \mathbb{P}[R_t \leq r] \\ &= \bar{Z}_1 \delta(r) + \left(\bar{Z}_1 + \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \right) \frac{1 - \rho_R \rho_E}{\rho_R \beta} \\ &\quad \cdot \exp \left[\frac{1 - \rho_E \rho_R}{\rho_R \beta} r \right] \mathbf{1}_{\{0 \leq r \leq R_{\max}\}} + \frac{1}{1 - \rho_E \rho_R} \\ &\quad \cdot \left(1 - (\rho_R \rho_E + \bar{Z}_1 1 - \rho_E \rho_R) \right. \\ &\quad \left. \cdot \exp \left[\frac{1 - \rho_E \rho_R}{\rho_R \beta} R_{\max} \right] \right) \delta(r - R_{\max}), \quad (26) \end{aligned}$$

where $\delta(\cdot)$ denotes the Dirac-delta function.

PROOF. Here, we show a condensed version of the proof by omitting various algebraic steps. See §6 in the e-companion to this paper for a detailed proof. From (2) and (16), we can show that

$$\mathbb{P}[R_{t+1} = 0 \mid R_t] = \mathbb{P}[R_{t+1} = 0] = \bar{Z}_1$$

and

$$\mathbb{P}[R_{t+1} = R_{\max} \mid R_t] = 1 - \bar{Z}_1 - \frac{R_{\max}}{\rho_R \beta} + \frac{(1 - \rho_E \rho_R) R_t}{\rho_R \beta}$$

in the steady-state. Also from (2), we can show that

$$f_{R_{t+1}|R_t}(u \mid R_t) = \begin{cases} \frac{\rho_E}{\beta} & \text{if } 0 < u < R_t \\ \frac{1}{\rho_R \beta} & \text{if } R_t \leq u < R_{\max}. \end{cases}$$

Therefore, we can write the conditional probability density function as

$$\begin{aligned} f_{R_{t+1}|R_t}(u \mid R_t = r) &= \bar{Z}_1 \delta(u) + \frac{\rho_E}{\beta} \mathbf{1}_{\{0 \leq u < r\}} + \frac{1}{\rho_R \beta} \mathbf{1}_{\{r \leq u \leq R_{\max}\}} \\ &\quad + \left(1 - \bar{Z}_1 - \frac{R_{\max}}{\rho_R \beta} + \frac{(1 - \rho_E \rho_R) r}{\rho_R \beta} \right) \delta(u - R_{\max}), \end{aligned}$$

where $\delta(\cdot)$ denotes the Dirac-delta function. Because

$$\mathbb{P}[R_t = 0] = \bar{Z}_1$$

in the steady state, we know that the stationary distribution can be written as

$$\begin{aligned} f_{R_t}(r) &= \bar{Z}_1 \delta(r) + g(r) \mathbf{1}_{\{0 \leq r \leq R_{\max}\}} \\ &\quad + \left(1 - \bar{Z}_1 - \int_{r=0}^{R_{\max}} g(r) dr \right) \delta(r - R_{\max}), \end{aligned}$$

for some function $g(r)$. By definition, the stationary distribution must satisfy

$$\begin{aligned} f_{R_{t+1}}(u) &= \int_{r=0}^{R_{\max}} f_{R_{t+1}|R_t}(u, r) dr \\ &= \int_{r=0}^{R_{\max}} f_{R_{t+1}|R_t}(u \mid R_t = r) f_{R_t}(r) dr = f_{R_t}(u). \end{aligned}$$

By computing the integral and matching the terms, we can show that

$$g(u) = \frac{\bar{Z}_1(1 - \rho_R \rho_E)}{\rho_R \beta} + \frac{\rho_R \rho_E}{\rho_R \beta} + \frac{1 - \rho_R \rho_E}{\rho_R \beta} \int_{r=0}^u g(r) dr.$$

Taking the derivative with respect to u on both sides gives

$$g'(u) = \frac{1 - \rho_R \rho_E}{\rho_R \beta} g(u).$$

Then, we can show that

$$g(r) = \left(\bar{Z}_1 + \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \right) \frac{1 - \rho_R \rho_E}{\rho_R \beta} \exp \left[\frac{1 - \rho_E \rho_R}{\rho_R \beta} r \right]$$

and

$$\begin{aligned} 1 - \bar{Z}_1 - \int_{r=0}^{R_{\max}} g(r) dr &= \frac{1}{1 - \rho_E \rho_R} \left(1 - (\rho_R \rho_E + \bar{Z}_1 1 - \rho_E \rho_R) \right. \\ &\quad \left. \cdot \exp \left[\frac{1 - \rho_E \rho_R}{\rho_R \beta} R_{\max} \right] \right). \quad \square \end{aligned}$$

The stationary distribution (26) shows that if the round-trip efficiency is lower, the probability of hitting the capacity limit R_{\max} is lower while the probability of depleting the storage is higher, as expected.

4.3. Economic Value of the Storage

Using the stationary distribution of the storage level obtained in the previous section, we can compute the following:

COROLLARY 2. *In steady state,*

$$\mathbb{E}[R_t] = \frac{R_{\max}}{1 - \rho_R \rho_E} - \left(\bar{Z}_1 + \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \right) \frac{\rho_R \beta}{1 - \rho_R \rho_E} \cdot \left(\exp \left[\frac{1 - \rho_R \rho_E}{\rho_R \beta} R_{\max} \right] - 1 \right) \quad (27)$$

and the expected revenue in steady state is

$$\begin{aligned} C_{SS}^{R_{\max}} &:= \mathbb{E}[C(S_t, x_t^*)] = \mu_p \rho_E \mathbb{E}[R_t] + \mu_p \mathbb{E}[\theta_t] \\ &\quad + \mu_p \bar{Z}_1 \beta - (m \mu_p + b) \frac{\beta}{2} \bar{Z}_2 \\ &= \frac{\mu_p \rho_E R_{\max}}{1 - \rho_R \rho_E} - \mu_p \beta \left(\bar{Z}_1 + \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \right) \frac{\rho_E \rho_R}{1 - \rho_R \rho_E} \\ &\quad \cdot \left(\exp \left[\frac{1 - \rho_R \rho_E}{\rho_R \beta} R_{\max} \right] - 1 \right) + \mu_p \beta \\ &\quad \cdot \left(\bar{Z}_1 - m \frac{\bar{Z}_2}{2} - \frac{1}{2} \right) + \mu_p \mu_Y - b \beta \frac{\bar{Z}_2}{2}. \end{aligned} \quad (28)$$

See §7 in the e-companion to this paper for the derivation of (27) and (28). We know that $K_1 = K_2 = 1$ if $R_{\max} = 0$. Therefore, from (28), if we do not have a storage and $R_{\max} = R_t = 0, \forall t$, the expected revenue at steady-state would be

$$C_{SS}^0 := \mu_p \beta \left(\bar{Z}_1 - m \frac{\bar{Z}_2}{2} - \frac{1}{2} \right) + \mu_p \mu_Y - b \beta \frac{\bar{Z}_2}{2}.$$

Then, the relative increase in the expected revenue in steady state due to the existence of storage is

$$\begin{aligned} \psi &:= \frac{C_{SS}^{R_{\max}} - C_{SS}^0}{C_{SS}^0} \\ &= \left\{ \frac{\rho_E \rho_R}{1 - \rho_R \rho_E} \frac{R_{\max}}{\beta \rho_R} - \left(\bar{Z}_1 + \frac{\rho_R \rho_E}{1 - \rho_R \rho_E} \right) \cdot \frac{\rho_E \rho_R}{1 - \rho_R \rho_E} \left(\exp \left[(1 - \rho_R \rho_E) \frac{R_{\max}}{\rho_R \beta} \right] - 1 \right) \right. \\ &\quad \left. + (\bar{Z}_1 - \bar{Z}_1) - \frac{1}{2} \left(m + \frac{b}{\mu_p} \right) (\bar{Z}_2 - \bar{Z}_2) \right\} \\ &\quad \cdot \left\{ \bar{Z}_1 - m \frac{\bar{Z}_2}{2} - \frac{1}{2} + \frac{\mu_Y}{\beta} - \frac{b}{\mu_p} \frac{\bar{Z}_2}{2} \right\}^{-1}. \end{aligned} \quad (29)$$

5. Numerical Results

In the previous section, we derived the optimal commitment policy and the corresponding value of the storage, assuming that the forecast of electricity generated from the wind farm is uniformly distributed. However, the hourly wind speed

data obtained from the North American Land Data Assimilation System (NLDAS) project shows that when forecasting the cube of the speed of the wind, a truncated Gaussian distribution fits the data better than a uniform distribution. In this section, we simulate the wind energy process using a truncated Gaussian distribution and compare the relative increase in revenue due to the existence of storage computed numerically by implementing our policy (16) to the one computed theoretically from Equation (29).

From the NLDAS project, we extracted wind speed data from 22 locations across the United States. Because the wind characteristics vary throughout the year due to seasonal effects, it is common to assume that the wind process is time invariant over a one-month period, but not beyond that (Ettoumi et al. 2003). Therefore, we use separate model parameters and corresponding policies for each month. We found that the third-order correlation is very small compared to the first- and second-order correlations, and represent $(Y_t)_{t \geq 1}$, the energy generated from our wind farm, as a second-order AR process:

$$Y_{t+1} = \mu_Y + \alpha_0(Y_t - \mu_Y) + \alpha_1(Y_{t-1} - \mu_Y) + \hat{y}_{t+1}, \quad (30)$$

for some μ_Y, α_0 , and α_1 . When we implement our policy (16), we assume $(\hat{y}_t)_{t \geq 1}$ is i.i.d. with distribution $\mathcal{U}(-\beta/2, \beta/2)$, for some β . β is computed by matching $\beta^2/12$ to the variance of the residual in the AR process, $(\hat{y}_t)_{t > 1}$. μ_Y s (in m^3/s^3) for the selected 22 locations computed using the January 2000 data, for example, are given in Table 1 and β s (in m^3/s^3) are given in Table 2. From Tables 1 and 2, we can see that μ_Y s and β s are comparable in magnitude, implying that wind energy production is highly volatile.

After we compute $\mu_Y, \alpha_0, \alpha_1$, and β using the NLDAS data, we generate wind energy processes $(Y_t)_{t \geq 1}$ from Equation (30) where $(\hat{y}_t)_{t > 1}$ is i.i.d. and $\hat{y}_t \sim \mathcal{N}(0, \beta^2/12), \forall t$. However, when we are computing our commitment from our policy (16), we assume $\hat{y}_t \sim \mathcal{U}(-\beta/2, \beta/2)$. From (6),

$$\theta_t = \mu_Y + \alpha_0(Y_t - \mu_Y) + \alpha_1(Y_{t-1} - \mu_Y) - \frac{\beta}{2}.$$

Next, we fit the hourly spot market price provided by a utility company to the process

$$p_{t+1} = \mu_p + (1 - \kappa \Delta \tau)(p_t - \mu_p) + \hat{p}_{t+1}, \quad (31)$$

where $(\hat{p}_t)_{t \geq 0}$ are i.i.d. with distribution $\mathcal{N}(0, \sigma_p^2)$. In our experiments, we use $\rho_E \rho_R = 0.75, \gamma = 0.99, \mu_p = 49.9, \sigma_p = 47.46, \Delta \tau = 1, \kappa = 0.4182, m = 1.6, b = 67.5$, and $(R_{\max})/(\rho_R \beta) = 0.5$. Then, $\bar{Z}_1 = 0.1912, \bar{Z}_2 = 0.0467, \tilde{Z}_1 = 0.3270$, and $\tilde{Z}_2 = 0.1139$. This gives

$$\psi = \frac{0.1893}{\mu_Y/\beta - 0.3411}. \quad (32)$$

We implemented our policy (16) 100 times by generating the prices from (31) and the wind energy process from (30)

Table 1. Mean of the cube of the speed of the wind in January 2000.

	120.0725 °W	111.3225 °W	102.5725 °W	93.8225 °W	85.0725 °W	76.3225 °W
51.8125 °N	181.7084	132.0368	144.4341	172.7166	276.2300	351.6345
46.1875 °N	173.3216	119.7605	318.7690	347.1192	482.2868	531.8000
40.5625 °N	156.4102	231.1095	N/A	380.6635	401.7359	491.8443
34.9375 °N	121.7831	N/A	224.7323	212.8919	198.2480	728.0436

Table 2. Spread of the cube of the speed of the wind in January 2000.

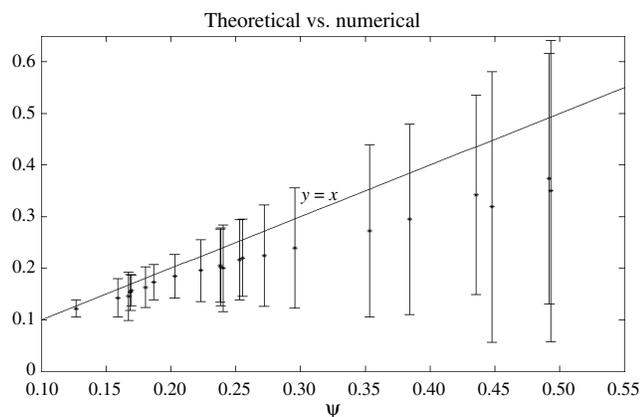
	120.0725 °W	111.3225 °W	102.5725 °W	93.8225 °W	85.0725 °W	76.3225 °W
51.8125 °N	250.3154	103.6640	186.1639	127.5475	150.6440	241.4260
46.1875 °N	159.1172	86.1458	294.4335	305.5882	329.9655	447.0356
40.5625 °N	150.7949	151.0185	N/A	456.5994	354.0033	501.0571
34.9375 °N	167.9374	N/A	294.1367	242.8136	175.6343	494.2871

using the coefficients μ , α_0 , α_1 , and β . Then, we computed the relative increase in revenue due to the existence of storage for each implementation of our policy and found the average of those values over the 100 experiments. Next, we computed the relative increase in revenue directly from Equation (29). The relative increase in revenue computed by implementing our policy (16), averaged over 36 months, is given in Table 3. The relative increase in revenue computed from (29), and hence (32), is given in Table 4.

From the above tables, we can see that the relative increase in revenue obtained through a sample run implementing our policy (16) is comparable to the relative increase in revenue computed using the closed-form Equation (29), even though the wind energy processes are actually generated from a truncated Gaussian distribution.

Figure 1 shows the relationship between the numerical results from Table 3 and the theoretical results from Table 4. There are 22 data points corresponding to each of the 22 locations. For each data point, the x -coordinate corresponds to the theoretical value computed from (29) and the y -coordinate corresponds to the numerical value computed from our policy (16). The error bar covers two standard deviations. In Figure 1, one can see that almost all

Figure 1. Plot of theoretical value of storage computed from (16) vs. value of storage computed numerically.



of the data points are slightly below the line $y = x$. That is, the relative increase in revenue computed from our policy (16) is almost always slightly less than the relative increase in revenue computed from (29). This is because the theoretical values were computed assuming that the wind energy process is generated from a uniform distribution, which makes our policy optimal, while the actual experiment used wind energy processes generated from a truncated Gaussian processes, making our policy suboptimal. The difference is approximately 15.6% on average.

6. Conclusions

In this paper, we have derived an optimal policy for making advance commitments of energy from an intermittent source such as wind, in the presence of a finite storage buffer, energy conversion losses, and a mean-reverting process for electricity prices. The goal of the paper was an analytical result that could be easily applied by energy economists, or as a heuristic within a simulation-based model. For this reason, we studied a stylized model that introduced several simplifying assumptions to make the problem analytically tractable. In addition to deriving an optimal policy for our stylized model, we were also able to derive an expression for the value of storage, making it possible to understand the interaction of volatility in wind, the capacity of the storage device, and storage losses.

Our model requires a number of assumptions such as stationarity in the wind and price processes, and the assumption of uniformly distributed errors in the wind forecast. It would be nice if we could show that the optimal policy always has a form similar to the newsvendor problem as shown in (16), regardless of the distribution of wind energy. Another dimension arises in risk mitigation when modeling heavy-tailed behaviors in electricity prices.

If the model were to be applied in the context of making day-ahead commitments, we have ignored the ability to make adjustments in the hour-ahead market. An important extension would be the derivation of a policy that captured the hour-ahead market within the day-ahead market.

Real-world energy storage tends to exhibit more complex physics than are assumed in simple inventory models. For

Table 3. Relative increase in revenue computed by implementing our policy (16).

	120.0725 °W	111.3225 °W	102.5725 °W	93.8225 °W	85.0725 °W	76.3225 °W
51.8125 °N	0.3682	0.1813	0.3378	0.1727	0.1213	0.1584
46.1875 °N	0.2185	0.1630	0.2191	0.2073	0.1538	0.1969
40.5625 °N	0.2245	0.1436	N/A	0.2985	0.2028	0.2394
34.9375 °N	0.3457	N/A	0.3211	0.2722	0.2005	0.1444

Table 4. Relative increase in revenue computed from (29).

	120.0725 °W	111.3225 °W	102.5725 °W	93.8225 °W	85.0725 °W	76.3225 °W
51.8125 °N	0.4919	0.2030	0.4356	0.1869	0.1268	0.1697
46.1875 °N	0.2530	0.1804	0.2553	0.2382	0.1689	0.2231
40.5625 °N	0.2719	0.1592	N/A	0.3843	0.2385	0.2955
34.9375 °N	0.4929	N/A	0.4476	0.3534	0.2403	0.1673

example, storage losses can be a function of the rate of energy production (which varies with the cube of the wind speed), and the amount of energy that can be stored in some devices can depend on the rate at which the energy has been stored. Finally, some devices such as compressed air require increasing amounts of energy as the device gets close to capacity (compressed air storage devices can reach pressures of 3,000 psi or more).

It is unlikely that we can derive analytical solutions for more general problems (for example, those that capture nonstationarities in the wind or price processes), but it is possible that a numerical solution could be used to calibrate an analytical model such as ours to reduce the errors due to these effects. Ultimately, there will always be a need for accurate models that will have to be solved using numerical methods, but at the same time we feel that there will also be interest in analytical models that are easy to compute and that provide insights into trade-offs between parameters.

It is possible that some of the issues that arise in the analysis of energy problems may spark new interest in problems in classical inventory theory that may share similar properties. For example, there are applications in classical supply chain management where the *supply* of product is random, and where vendors may have to make commitments to deliver product, using stored inventories to help smooth over supply problems.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Acknowledgments

The authors thank Eric Wood of the Civil and Environmental Engineering Department at Princeton University for providing the NLDAS wind speed data. The authors would also like to thank PJM and NRG Energy for providing information regarding the electricity market system. The research was supported in part by SAP and the Air Force Office of Scientific Research, AFOSR contract FA9550-08-1-0195.

References

- Axsäter, S. 2000. *Inventory Control*. Kluwer Academic Publishers, Boston.
- Brown, P., J. Matos. 2008. Optimization of pumped storage capacity in an isolated power system with large renewable penetration. *IEEE Trans. Power Systems* **23**(2) 523–531.
- Brunetto, C., G. Tina. 2007. Optimal hydrogen storage sizing for wind power plants in day ahead electricity market. *Renewable Power Generation* **1**(4) 220–226.
- Castronuovo, E., J. Lopes. 2004. On the optimization of the daily operation of a wind-hydro power plant. *IEEE Trans. Power Systems* **19**(3) 1599–1606.
- Cosgrove, B. A., D. Lohman, K. E. Mitchell, P. R. Houser, E. F. Wood, J. C. Schaake, A. Robock. 2003. Real-time and retrospective forcing in the North American land data assimilation system (NLDAS) project. *J. Geophysical Res.* **108**(22) 8842–8853.
- Ettoumi, F. Y., H. Sauvageot, A. E. H. Adane. 2003. Statistical bivariate modelling of wind using first-order Markov chain and Weibull distribution. *Renewable Energy* **28**(11) 1787–1802.
- Eydeland, A., K. Wolyniec. 2003. *Energy and Power Risk Management*. John Wiley & Sons, Hoboken, NJ.
- García-González, J., R. M. R. de la Muela, L. M. Santos, A. M. González. 2008. Stochastic joint optimization of wind generation and pumped storage units in an electricity market. *IEEE Trans. Power Systems* **23**(2) 460–468.
- Ibrahim, H., A. Ilinca, J. Perron. 2008. Energy storage systems—characteristics and comparisons. *Renewable and Sustainable Energy Rev.* **12**(5) 1221–1250.
- Khouja, M. 1999. The single-period newsvendor problem: Literature review and suggestions for future research. *Omega* **27**(5) 537–553.
- Korpaas, M., A. T. Holen, R. Hildrum. 2003. Operation and sizing of energy storage for wind power plants in a market system. *Internat. J. Electr. Power and Energy Systems* **25**(8) 599–606.
- MacKerron, G., P. Pearson, eds. 2000. *International Energy Experience: Markets, Regulation and the Environment*. Imperial College Press, London.
- Morthorst, P. E. 2003. Wind power and the conditions at a liberalized power market. *Wind Energy* **6**(3) 297–308.
- Nandalal, K. D. W., J. J. Bogardi. 2007. *Dyanamic Programming Based Operation of Reservoirs: Applicability and Limits*. Cambridge University Press, New York.
- Petruzzi, N., M. Dada. 1999. Pricing and the newsvendor problem. *Oper. Res.* **40**(5) 582–596.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York.
- Sioshansi, R., P. Denholm, T. Jenkin, J. Weiss. 2009. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Econom.* **31**(2) 269–277.
- Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill, Boston.